Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Integrating Artificial Intelligence into SAP ERP: Opportunities, Challenges, and Implementation Paths

Jim Dorsey

jimdorsey4521@gmail.com

Miami University of Science and Technology

Research Article

Keywords: SAP ERP, Artificial Intelligence, Clean Core, Enterprise Resource Planning, Implementation Paths, AI Governance

Posted Date: October 8th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11284801/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

The integration of Artificial Intelligence into SAP ERP systems has emerged as a defining challenge in enterprise digital transformation. This paper systematically examines the opportunities, challenges, and implementation paths of AI integration within SAP ERP environments. In terms of opportunities, SAP's Business AI platform offers a multi-layered capability framework ranging from embedded functionalities to custom intelligent agents, covering key scenarios such as demand forecasting, process automation, and predictive maintenance. However, implementation faces three structural challenges: the tension between Clean Core principles and AI customization requirements, the context gap arising from the absence of enterprise-specific business context, and the inadequate adaptation of AI ethics and governance frameworks to ERP environments. Regarding implementation paths, both academic and practitioner consensus leans toward a dual-track strategy: adopting sidecar custom AI in traditional ECC environments to protect core stability, while progressively activating embedded AI capabilities in cloud ERP environments, both contingent upon data readiness as a prerequisite. This paper argues that successful AI integration depends not on technological sophistication, but on establishing a sustainable balance between the stability requirements of ERP and the adaptive demands of AI.

1. Introduction

1.1 Research Background

The convergence of enterprise resource planning and artificial intelligence has become one of the most consequential developments in enterprise information systems over the past decade. ERP systems serve as the transactional backbone for organizations across manufacturing, finance, healthcare, and public sectors, integrating disparate functional areas into unified data architectures that support operational visibility and strategic decision-making (Katuu, 2020). As these systems have matured, the question of how to augment their capabilities with artificial intelligence has shifted from speculative inquiry to urgent practical concern.

SAP occupies a distinctive position in this landscape. As the dominant ERP vendor globally, with a customer base exceeding 400,000 organizations, SAP's strategic direction shapes the technological horizons of a substantial portion of the world's enterprise computing infrastructure (SAP SE, 2023). The company's pivot toward intelligent enterprise solutions represents a fundamental reorientation of its product philosophy, moving from systems of record toward systems of intelligence that embed AI capabilities directly into core business processes (SAP SE, 2023). This transition is not merely additive; it requires rethinking the architectural, governance, and organizational foundations upon which ERP value has traditionally been built.

The timing of this transition is inseparable from the operational realities facing SAP customers. Mainstream maintenance for SAP ECC 6.0, the predecessor system upon which many organizations continue to rely, ends in 2027, with extended maintenance available only through 2030 at premium cost

(SAP SE, 2023). This deadline creates a dual imperative: organizations must migrate to SAP S/4HANA while simultaneously determining how to integrate AI capabilities that are increasingly built exclusively for cloud infrastructure. The migration and the AI integration are not sequential projects but intertwined challenges that must be addressed within a compressed timeframe.

1.2 Problem Statement

Despite growing interest in AI-enabled ERP systems, a significant gap persists between aspiration and implementation. Existing research has examined specific AI applications within ERP contexts, including machine learning for demand forecasting, natural language processing for user interaction, and predictive analytics for maintenance scheduling (Almeida et al., 2020). However, these studies tend to focus on isolated use cases or theoretical frameworks rather than providing comprehensive guidance on how organizations can systematically integrate AI capabilities while preserving core system stability.

A particular tension characterizes the contemporary discourse. On one hand, SAP's architecture philosophy emphasises Clean Core principles: the maintenance of standardised, upgrade-stable systems through the minimisation of custom modifications (SAP SE, 2022). On the other hand, AI implementations frequently require access to organisation-specific data, bespoke logic, and adaptation to unique business processes. This tension is not merely technical but organisational, raising questions about who bears responsibility for AI lifecycle management, how governance structures should be designed, and what capabilities organisations must develop internally versus acquiring from vendors.

The context gap further complicates integration. Vendor-delivered AI scenarios address standardised, widely recurring processes, but organisations require enrichment with proprietary data fields and workflows that reflect their distinctive competitive positioning (SAP SE, 2023). When AI modules are deployed without adequate attention to this contextual dimension, they risk becoming functionally inert, what some observers term ghost technology: systems that are technically operational but produce no meaningful organisational value because they were not integrated into the actual work practices they were intended to support.

1.3 Research Questions

This paper addresses three interrelated questions:

First, how can organisations systematically integrate AI capabilities into SAP ERP environments while preserving core system stability and upgrade readiness?

Second, what structural challenges distinguish AI-ERP integration from traditional ERP customisation, and how do these challenges manifest across different organisational contexts?

Third, what implementation frameworks and governance mechanisms enable sustainable AI adoption within the constraints imposed by Clean Core principles and the ongoing ECC-to-S/4HANA transition?

1.4 Significance and Contribution

This research makes both theoretical and practical contributions. Theoretically, it synthesises fragmented literature on AI-ERP integration into a coherent framework that addresses the lifecycle dimensions of implementation, from initial assessment through governance and continuous improvement. It advances understanding of the socio-technical dynamics that shape AI adoption in enterprise systems, moving beyond purely technical or purely organisational accounts to examine their interaction.

Practically, the paper provides guidance for organisations navigating the dual-track reality of ECC and S/4HANA environments. It delineates specific extension patterns that preserve Clean Core compliance while enabling AI innovation, and it proposes a phased roadmap that reflects the staged nature of enterprise AI adoption. The framework acknowledges that organisations operate under resource constraints, regulatory requirements, and legacy system dependencies that cannot be assumed away.

1.5 Scope and Boundaries

This study focuses on the SAP-specific ecosystem, including S/4HANA (on-premise, private cloud, and public cloud editions), the Business Technology Platform, and the emerging Business AI Platform. While recognising that organisations may operate heterogeneous ERP landscapes involving multiple vendors, the SAP context provides sufficient coherence for meaningful analysis given the company's distinctive architectural philosophy and market position.

The emphasis is on organisational and methodological dimensions rather than low-level technical implementation details. Continuous integration and deployment pipelines, infrastructure provisioning, and model training at the code level are beyond the scope of this analysis, except where they intersect with governance and architecture decisions. The paper acknowledges regional and regulatory variation without attempting exhaustive coverage of every jurisdiction-specific requirement.

2. Literature Review and Theoretical Foundations

2.1 ERP Systems: Evolution and Architecture

Enterprise resource planning systems represent the culmination of several decades of evolution in enterprise computing, from material requirements planning in the 1960s and 1970s, through manufacturing resource planning in the 1980s, to the integrated ERP architectures that emerged in the 1990s (Katuu, 2020). These systems promised a single source of truth, integrating finance, supply chain, human resources, and manufacturing into unified databases and process flows.

SAP's architectural evolution reflects broader industry shifts. The company's flagship product, SAP ECC, dominated the ERP market for nearly two decades, becoming the de facto standard for large enterprises. Its successor, S/4HANA, introduced a fundamentally different data model based on in-memory computing, enabling real-time analytics and simplified data structures (SAP SE, 2022). More recently,

SAP has positioned itself around the concept of the intelligent enterprise, integrating AI, analytics, and cloud infrastructure into a unified platform vision (SAP SE, 2023).

The Clean Core concept has emerged as a central architectural principle in SAP's contemporary guidance. Clean Core refers to the maintenance of standardised, unmodified SAP code, with customisations and extensions channelled through published APIs and platform services rather than direct modifications to core objects (SAP SE, 2022). The rationale is straightforward: modifications to core code complicate upgrades, introduce technical debt, and create dependencies that constrain future flexibility. For AI integration, Clean Core has profound implications because it constrains how custom models and organisation-specific logic can be incorporated into the ERP environment.

2.2 Artificial Intelligence in Enterprise Contexts

Artificial intelligence in enterprise contexts encompasses a diverse set of technologies, including machine learning, natural language processing, computer vision, and, more recently, generative AI. Each of these capability classes has distinct characteristics that shape its applicability to ERP scenarios (Almeida et al., 2020). Machine learning excels at pattern recognition and prediction, making it suitable for demand forecasting, anomaly detection, and predictive maintenance. Natural language processing enables conversational interfaces and document processing. Generative AI, while newer in enterprise deployment, offers capabilities in content creation, code generation, and complex reasoning tasks.

A key architectural distinction separates embedded AI from sidecar AI. Embedded AI refers to capabilities built directly into the ERP application, leveraging the system's native data and process context. Sidecar AI, by contrast, operates alongside the ERP system, consuming data through APIs and returning results that the ERP system then processes (SAP SE, 2023). Each approach has trade-offs: embedded AI offers seamless integration and lower latency but limited customisation; sidecar AI offers flexibility and independence but introduces complexity in data synchronisation and governance.

The concept of agentic AI has gained prominence in recent discourse, referring to AI systems that can autonomously execute multi-step tasks, make decisions, and orchestrate complex workflows (Russell & Norvig, 2021). In ERP contexts, agentic AI raises significant governance questions, particularly regarding accountability and the boundaries of autonomous action.

2.3 Prior Research on AI-ERP Integration

The academic literature on AI-ERP integration has expanded considerably over the past five years, though significant gaps remain. A systematic review of the field reveals several patterns (Almeida et al., 2020). The dominant methodologies employed in AI-ERP research are machine learning, deep learning, and natural language processing, with generative AI representing an emerging but still limited area of investigation. Application domains are unevenly distributed, with finance and supply chain receiving disproportionate attention, while production planning and quality management remain comparatively understudied.

Documented benefits of AI-ERP integration include enhanced decision-making through predictive analytics, operational efficiency through process automation, cost reduction through improved resource allocation, and improved customer experience through personalisation (Almeida et al., 2020). These benefits are frequently cited in vendor literature and case studies, though independent empirical validation remains limited.

Documented challenges cluster around several themes. Data quality and availability represent persistent obstacles, as AI models require clean, structured, and sufficiently voluminous data to function effectively. Integration complexity arises from the need to connect AI services with legacy ERP architectures. Security and privacy concerns are amplified when AI systems process sensitive business data. Regulatory compliance requirements, particularly in regulated industries and jurisdictions, impose additional constraints. User competence and organisational readiness are frequently cited as underestimated factors in implementation success (Almeida et al., 2020).

2.4 Identified Research Gaps

Despite this growing body of work, several gaps persist. First, existing research lacks lifecycle-oriented frameworks for AI-ERP integration, particularly frameworks that account for the DevOps practices required to manage AI models over time. Most studies focus on initial implementation rather than the ongoing monitoring, retraining, and governance that AI systems require.

Second, the socio-technical dimensions of AI integration are insufficiently addressed. ERP implementations have long been recognised as organisational change projects as much as technical projects (Katuu, 2020), yet AI integration research often treats organisational factors as secondary considerations rather than constitutive elements of success.

Third, there is limited empirical evaluation of comprehensive implementation methodologies. Many studies propose frameworks but do not test them in practice, leaving questions about their applicability across different organisational contexts.

Fourth, certain functional domains remain underrepresented. The concentration of research on finance and supply chain applications leaves production planning, quality management, and human capital management comparatively unexplored, despite their importance in many organisations.

These gaps motivate the present study, which seeks to synthesise existing knowledge into a coherent framework while addressing the lifecycle and socio-technical dimensions that prior work has tended to neglect.

3. Research Methodology

3.1 Research Design

This study adopts a qualitative, exploratory approach combining systematic literature review with conceptual framework development. The choice of a non-empirical design reflects several considerations. The field of AI-ERP integration is emergent and rapidly evolving, making large-scale empirical studies difficult to execute within meaningful timeframes. Vendor product portfolios change frequently, and the organisational practices surrounding AI adoption are still in formative stages. Under these conditions, a systematic synthesis of existing knowledge, coupled with conceptual framework development, offers greater value than premature empirical testing.

The research design follows the logic of conceptual research as articulated in information systems scholarship, where the researcher synthesises disparate literature streams into coherent theoretical frameworks that can guide subsequent empirical investigation (Webster & Watson, 2002). The framework developed here is intended to be actionable, providing specific guidance for practitioners while remaining theoretically grounded.

3.2 Systematic Literature Review Protocol

The literature review followed a structured protocol informed by established guidelines for systematic reviews in information systems and management research (Webster & Watson, 2002). Database selection included Scopus, IEEE Xplore, ScienceDirect, and the AIS eLibrary, covering both technical and managerial literatures. Search terms combined variations of ERP and SAP with AI-related concepts, including machine learning, artificial intelligence, and intelligent systems.

Inclusion criteria specified peer-reviewed journal articles, conference proceedings, and book chapters published in English. Exclusion criteria eliminated purely technical papers focused on algorithm development without organisational context, vendor marketing materials, and studies where AI was mentioned only peripherally. The screening process yielded a corpus of studies that were then analysed thematically to identify recurring patterns, contradictions, and gaps.

Thematic analysis proceeded through iterative coding, with initial codes clustered into broader themes that structure the literature review and inform the framework development. This approach is consistent with established practices in qualitative synthesis (Webster & Watson, 2002).

3.3 Framework Development Approach

The conceptual framework developed in this paper synthesises existing DevOps and AI lifecycle models, adapting them to the specific constraints of ERP environments. Three considerations shape this adaptation. First, Clean Core principles impose architectural constraints that distinguish ERP AI integration from generic AI deployment. Second, the upgrade stability requirement in ERP environments creates unique challenges for model management and version control. Third, the delineation of vendor versus customer responsibilities in the SAP ecosystem determines which aspects of the AI lifecycle organisations can control directly and which they must influence through configuration and extension.

The framework is presented in Chap. 6 and comprises several components: a dual-track implementation strategy, a DevOps process model, specific extension patterns compatible with Clean Core principles, a phased roadmap, and a governance architecture.

3.4 Limitations

Several limitations should be acknowledged. The reliance on published literature and vendor documentation, rather than primary empirical data, means the framework developed here requires subsequent validation. The rapid pace of change in AI and ERP technologies creates obsolescence risk; specific product names and capabilities may evolve even as the underlying architectural principles remain stable.

The available literature exhibits geographic and sectoral bias. Studies from North American and European contexts predominate, and certain industries, notably manufacturing and finance, are overrepresented. While the framework is intended to be generalisable, its applicability to other contexts requires careful assessment.

Finally, the exclusion of certain technical dimensions, such as detailed model architecture decisions or infrastructure provisioning, means the framework operates at a level of abstraction that may not satisfy readers seeking implementation-level guidance. This is a deliberate choice, reflecting the paper's focus on organisational and methodological dimensions, but it does constitute a boundary of the contribution.

4. Opportunities: The AI-Enabled SAP Landscape

4.1 SAP Business AI: Architecture and Capability Layers

SAP's approach to AI integration is structured around a multi-layered architecture that spans infrastructure, platform, and application levels. At the foundation, the Business Technology Platform provides the runtime environment, data management capabilities, and integration services upon which AI applications depend (SAP SE, 2023). The Business Data Cloud layer provides governed access to enterprise data, addressing the critical requirement that AI models be trained and operated on reliable, contextualised information. The AI Foundation layer offers pre-trained models, development tools, and runtime services that organisations can leverage to build custom AI capabilities or deploy vendor-delivered scenarios.

This architecture embodies what SAP describes as an AI flywheel logic: applications generate data, data trains AI models, and improved AI enhances applications, creating a virtuous cycle that compounds value over time (SAP SE, 2023). The flywheel metaphor captures an important insight: AI value in ERP contexts is not a one-time implementation but an ongoing process of learning and refinement. Organisations that treat AI as a discrete project rather than a continuous capability are unlikely to realise the full benefits.

Responsible AI considerations are embedded as a cross-cutting governance layer. SAP has articulated principles around explainability, fairness, robustness, and human oversight that are intended to apply across all AI capabilities (SAP SE, 2023). These principles are operationalised through technical mechanisms such as model documentation, bias testing, and human-in-the-loop workflows, as well as through governance processes that define accountability and escalation paths.

4.2 Embedded AI Use Cases Across Business Domains

The opportunities for AI integration span functional domains, with varying degrees of maturity and documented impact.

In finance, AI capabilities address several high-volume, rules-intensive processes. Predictive invoice matching uses machine learning to identify likely matches between invoices and purchase orders, reducing manual reconciliation effort and accelerating payment cycles. Automated revenue recognition applies pattern recognition to complex contract structures, supporting compliance with accounting standards while reducing manual review burden. Trade compliance systems use AI to screen transactions against regulatory requirements, flagging potential violations for human review. Intelligent collections prioritise receivables based on payment likelihood, optimising cash flow management (SAP SE, 2023).

In supply chain management, demand forecasting represents one of the most mature AI applications in ERP contexts. Machine learning models incorporate historical sales data, promotional calendars, economic indicators, and external signals to generate forecasts that outperform traditional statistical methods. Predictive maintenance uses sensor data and operational history to anticipate equipment failures before they occur, reducing unplanned downtime. Autonomous warehousing combines computer vision, robotics, and AI-driven orchestration to optimise picking, packing, and inventory management (Almeida et al., 2020).

In customer experience, intelligent sales execution uses AI to prioritise leads, recommend products, and automate routine customer interactions. Service automation applies natural language processing to customer inquiries, routing them to appropriate resolution paths and, where appropriate, resolving them without human intervention. Marketing optimisation uses predictive analytics to personalise campaign targeting and content delivery.

In human capital management, workforce planning uses predictive models to anticipate staffing needs and identify skills gaps. Benefits support applies conversational AI to employee inquiries about compensation and benefits, reducing HR service centre load. Learning and development platforms use AI to recommend personalised learning paths based on role requirements and career aspirations.

In procurement, contract price optimisation uses market data and historical patterns to identify opportunities for cost reduction. Supplier risk monitoring applies predictive analytics to external signals and internal performance data to flag potential disruptions. Spend analytics uses natural language

processing to categorise and analyse procurement data, identifying consolidation opportunities and maverick spending (SAP SE, 2023).

4.3 The Maturity Progression: From Embedded AI to Agentic Autonomy

SAP's product roadmap envisions a progression from embedded AI through agentic AI to industry-specific autonomous scenarios. This progression is analytically useful for understanding the opportunities available at different stages of organisational maturity.

At the first level, embedded AI provides contextual guidance, search capabilities, and summarisation within existing workflows. Examples include intelligent search that surfaces relevant documents and transactions, summarisation of lengthy records or histories, and contextual recommendations embedded in process screens. These capabilities augment human decision-making without fundamentally altering process flows (SAP SE, 2023).

At the second level, agentic AI introduces multi-agent orchestration, where AI systems coordinate across process steps and system boundaries to execute complex workflows. Human checkpoints remain, but the locus of execution shifts from the human user to the AI system. Examples include automated resolution of routine customer inquiries, end-to-end processing of standard procurement transactions, and proactive identification and resolution of supply chain disruptions.

At the third level, industry AI offers sector-specific autonomous scenarios that combine deep domain knowledge with AI capabilities. These scenarios address processes unique to particular industries, such as clinical trial matching in life sciences, claims processing in insurance, or production scheduling in discrete manufacturing.

The conversational interface, exemplified by SAP's Joule, serves as a unifying layer across these maturity levels. Joule provides a natural language interface to the full range of AI capabilities, allowing users to interact with the system conversationally rather than navigating complex transaction screens. This interface paradigm has significant implications for user adoption and training requirements (SAP SE, 2023).

4.4 Quantifiable Opportunities

The literature reports several quantifiable outcomes from AI-ERP integration, though independent validation remains limited. SAP's internal deployment of AI capabilities across its own operations has been cited as demonstrating productivity gains, though the specific metrics are vendor-reported and should be interpreted with appropriate caution.

Case studies from early adopters describe process improvements ranging from reduced cycle times to improved forecast accuracy. These reports are consistent with broader literature on AI in enterprise contexts, which documents efficiency gains and improved decision quality across multiple domains

(Almeida et al., 2020). However, the literature also cautions that benefits are not automatic and depend on factors including data quality, user adoption, and process alignment.

The most credible opportunities may be those that are incremental and well-bounded: reducing manual effort in high-volume processes, improving forecast accuracy through better use of available data, and surfacing insights that would otherwise require extensive manual analysis. Transformational opportunities, while compelling in vendor narratives, require organisational capabilities that many enterprises have not yet developed.

5. Challenges: Structural Barriers to Integration

5.1 The Clean Core Tension

The most fundamental challenge in AI-ERP integration arises from the tension between Clean Core principles and the requirements of AI customisation. Clean Core principles hold that ERP systems should remain as close to standard as possible, with customisations channelled through released APIs and platform services rather than direct modifications to core objects (SAP SE, 2022). The rationale is compelling: modifications to core code complicate upgrades, introduce technical debt, and create dependencies that constrain future flexibility.

AI implementations, however, frequently require access to organisation-specific data, bespoke logic, and adaptation to unique business processes. A predictive maintenance model, for example, may need to incorporate sensor data from proprietary equipment, historical maintenance records with organisation-specific coding, and operational context that is not captured in standard SAP fields. A demand forecasting model may need to incorporate promotional calendars, competitor actions, and market intelligence that reside outside the ERP system.

The problem is compounded by the classification of SAP objects. SAP categorises objects as released, stable, or internal, with corresponding recommendations for extension. Released objects are intended for customer extension and carry compatibility guarantees. Internal objects are subject to change without notice and are not recommended for extension. Many AI use cases, however, require access to data or logic that resides in internal objects, creating a conflict between what is technically feasible and what is architecturally advisable (SAP SE, 2022).

The classic API versus released API distinction further complicates the landscape. Classic APIs, while technically accessible, may not carry the same stability guarantees as released APIs, creating upgrade risk for organisations that build AI capabilities upon them. This risk is not merely theoretical; organisations that have built custom extensions on classic APIs have experienced upgrade disruptions when SAP changed underlying object structures.

5.2 The Context Gap

A second structural challenge arises from what may be termed the context gap. Vendor-delivered AI scenarios necessarily address standardised, widely recurring processes, as these represent the common denominator across the customer base. Organisations, however, require enrichment with proprietary data fields, bespoke workflows, and organisation-specific business logic that reflect their distinctive competitive positioning (SAP SE, 2023).

The context gap manifests in several ways. First, AI models trained on generic data may underperform when applied to organisation-specific contexts. A demand forecasting model trained on aggregate industry patterns may not capture the idiosyncrasies of a particular company's customer base, promotional strategies, or supply chain relationships. Second, AI recommendations that ignore organisation-specific constraints may be impractical or even counterproductive. A cost optimisation recommendation that ignores contractual obligations or strategic supplier relationships may save money in the short term while damaging long-term value.

Third, the context gap affects user trust and adoption. When AI systems produce recommendations that are obviously inconsistent with organisational reality, users quickly learn to disregard them. This dynamic can render AI modules functionally inert even when they are technically operational, a phenomenon that has been observed in multiple enterprise AI deployments.

5.3 The Data Readiness Problem

AI depends on governed, structured, and reliable data, yet many organisations struggle with data quality issues that undermine AI effectiveness. Inconsistent master data, siloed information across business units, and reliance on manual data entry create conditions in which AI models cannot function as intended (Almeida et al., 2020).

The garbage in, garbage out dynamic is well understood in principle but frequently underestimated in practice. ERP systems accumulate data over years or decades, with varying standards of quality and consistency. Mergers and acquisitions introduce additional complexity, as organisations must reconcile data from different systems and different conventions. Even within a single organisation, data quality may vary significantly across business units and process areas.

For AI specifically, the data readiness problem is amplified by the probabilistic nature of AI outputs. Unlike deterministic systems, which produce consistent outputs given consistent inputs, AI models may produce varying outputs in response to minor data variations. This characteristic makes data quality even more critical, as small inconsistencies can produce significant variations in AI recommendations.

The data readiness problem is not merely technical. It reflects organisational issues including unclear data ownership, misaligned incentives, and insufficient investment in data governance. Addressing data readiness therefore requires organisational change as much as technical intervention.

5.4 Governance and Ethical Challenges

AI integration raises governance and ethical challenges that existing ERP governance frameworks are often ill-equipped to address. These challenges span several domains.

Bias in AI systems represents a significant concern, particularly in high-stakes decisions such as recruitment, resource allocation, and performance evaluation. AI models trained on historical data may perpetuate or amplify existing biases, producing recommendations that disadvantage certain groups. In ERP contexts, where AI may influence decisions about hiring, promotion, and resource distribution, the consequences of bias can be substantial.

Opacity of decision logic poses additional challenges. Many AI models, particularly deep learning models, operate as black boxes, producing recommendations without clear explanations of how those recommendations were derived. In ERP contexts, where decisions may need to be justified to auditors, regulators, or affected parties, this opacity creates accountability problems.

Privacy and data protection risks are amplified when AI systems process sensitive business data. Regulatory frameworks, including the General Data Protection Regulation in Europe and analogous regulations in other jurisdictions, impose requirements around data minimisation, purpose limitation, and individual rights that AI systems must respect. Compliance requires careful attention to data flows, model training practices, and the handling of AI outputs that contain personal data.

Accountability gaps arise when AI systems make or inform decisions. Who is responsible when an autonomous agent errs? The vendor that supplied the AI capability? The organisation that deployed it? The individual who configured it? The user who accepted its recommendation? These questions lack clear answers in many organisations, creating uncertainty that can delay AI adoption or lead to inadequate oversight.

Environmental costs of computational intensity represent an emerging concern. Large AI models require substantial computational resources for training and inference, with associated energy consumption and carbon emissions. Organisations with sustainability commitments must account for these costs, which may conflict with efficiency objectives.

Regulatory fragmentation adds complexity. Different jurisdictions impose different requirements, and sector-specific regulations may impose additional constraints. Organisations operating across multiple jurisdictions must navigate this complexity, often without clear guidance on how to reconcile conflicting requirements (SAP SE, 2023).

5.5 Organisational and Human Factors

The socio-technical complexity of ERP implementations has long been recognised (Katuu, 2020). AI integration adds new dimensions to this complexity, including shifts in power relations, changes to organisational routines, and new competence requirements.

Power relations are affected when AI systems redistribute decision-making authority. When AI recommends or decides, the locus of control shifts from human actors to algorithms and those who design them. This shift may be resisted by individuals who perceive their expertise or autonomy to be threatened. Understanding these dynamics is essential for successful implementation.

Organisational routines are disrupted when AI changes how work is performed. Tasks that were previously performed by humans may be automated, while new tasks emerge around AI supervision and exception handling. These changes require adaptation at both individual and organisational levels, and the process of adaptation is rarely smooth.

User competence and adoption barriers represent persistent challenges. AI systems may require new skills, including data literacy, prompt engineering, and the ability to interpret AI outputs critically. Organisations may lack these skills internally and may struggle to develop them quickly enough to keep pace with technological change.

The shift in workforce roles is significant and often underestimated. The transition from data entry to AI monitoring and steering requires different skills and different attitudes. Individuals who have built careers around executing transactions may find the shift to overseeing AI systems disorienting or threatening. Change management programmes must address these concerns directly.

The hollow feedback loop risk deserves particular attention. AI systems learn from data generated by existing processes, which may reflect existing institutional dysfunction. If the underlying processes are flawed, AI may amplify rather than correct those flaws. This dynamic has been observed in various domains, where AI systems perpetuate rather than resolve existing problems.

5.6 The ECC-to-S/4HANA Transition as Complicating Factor

The ongoing transition from SAP ECC to S/4HANA creates a particularly complex context for AI integration. Mainstream maintenance for ECC ends in 2027, with extended maintenance available only through 2030, creating a deadline that forces organisations to migrate (SAP SE, 2023). Yet SAP's AI innovations are built exclusively for cloud infrastructure, meaning that organisations cannot access the full range of AI capabilities until they have migrated.

The result is a dual-track reality. Organisations may operate ECC for core transactional processes while simultaneously running S/4HANA in selected areas, or they may maintain ECC while building AI capabilities on separate platforms. This dual-track approach introduces complexity in data integration, process coordination, and governance. It also creates uncertainty about where to invest in AI capabilities: on the current ECC platform, where capabilities may be limited, or on the future S/4HANA platform, where returns may not be realised for years.

The dual-track reality also affects the Clean Core calculus. Organisations that have customised their ECC systems extensively may face difficult choices about whether to carry those customisations forward to

S/4HANA. Customisations that were justified under ECC may not be compatible with S/4HANA's architecture, and the migration process may force organisations to rationalise their extensions. AI integration adds another dimension to this rationalisation, as organisations must determine which AI capabilities to embed in core processes and which to implement as sidecar solutions.

6. Implementation Paths: Frameworks and Strategies

6.1 The Dual-Track Strategy

The dual-track strategy emerges from the recognition that organisations operate in heterogeneous environments and cannot wait for full migration to begin realising AI value. The strategy involves two parallel tracks, each with distinct characteristics.

Track one addresses traditional ECC environments. In these environments, the priority is protecting core stability while enabling AI innovation. This is achieved through sidecar AI implementations, where AI capabilities run on separate platforms and interact with ECC through published APIs. This approach preserves Clean Core principles by avoiding modifications to ECC code while still enabling AI capabilities. The trade-offs include additional integration complexity and potential latency, but these are often acceptable given the alternative of destabilising core systems.

Track two addresses cloud ERP environments, particularly S/4HANA Cloud. In these environments, the priority is progressive activation of embedded AI capabilities. Organisations can begin with out-of-the-box AI scenarios that SAP delivers as part of the standard product, then extend these scenarios with custom models and organisation-specific logic as capabilities mature. The cloud environment provides the infrastructure and governance framework for AI lifecycle management, reducing the burden on customer IT organisations.

Both tracks are contingent upon data readiness. AI cannot function effectively on poor-quality data, and organisations must invest in data governance, master data management, and data quality improvement before or alongside AI implementation. This requirement is often underestimated, leading to disappointing outcomes when organisations attempt to deploy AI on data that is not fit for purpose.

6.2 A DevOps Framework for AI-ERP Integration

The literature on AI-ERP integration lacks lifecycle-oriented frameworks that address the ongoing management of AI capabilities (Almeida et al., 2020). To address this gap, a DevOps framework is proposed that adapts established practices from software development to the specific constraints of ERP environments.

The framework comprises five phases. The first phase, gap analysis, involves assessing the organisation's current state across dimensions including data quality, process standardisation, and AI competence, and comparing this state to the requirements of intended AI use cases. The output is a prioritised list of gaps and a plan for addressing them.

The second phase, requirements derivation, involves translating business objectives into technical requirements for AI capabilities. This phase must account for the constraints imposed by Clean Core principles and the available extension mechanisms.

The third phase, framework design, involves selecting the appropriate architecture for AI integration. Options include embedded AI, sidecar AI, or hybrid approaches, each with distinct implications for Clean Core compliance, data flows, and governance.

The fourth phase, implementation, involves building, testing, and deploying the AI capability. This phase follows established DevOps practices including version control, automated testing, and continuous integration, adapted to the specifics of AI models and ERP integration.

The fifth phase, evaluation, involves monitoring AI performance against defined metrics, identifying degradation or drift, and initiating retraining or adjustment as needed. This phase is continuous rather than terminal, reflecting the ongoing nature of AI lifecycle management.

A critical feature of the framework is the delineation of vendor versus customer responsibilities. The vendor provides technical capability, including APIs, pre-trained models, and embedded analytics. The customer must manage the full AI lifecycle, including data preparation, model customisation, integration, and ongoing monitoring. Failure to recognise this division of responsibility leads to unrealistic expectations and inadequate investment.

6.3 Clean Core-Compatible Extension Patterns

The tension between Clean Core principles and AI customisation requirements can be managed through appropriate selection of extension patterns. Four patterns merit consideration, ordered from most to least preferred from a Clean Core perspective.

Side-by-side extensibility on SAP BTP represents the primary recommendation for AI workloads requiring custom models or organisation-specific logic. In this pattern, the AI capability runs on the Business Technology Platform, consuming data from SAP systems through released APIs and returning results that the core system processes. This pattern preserves Clean Core compliance because no modifications are made to core objects. It offers flexibility and independence, though it introduces integration complexity and potential latency.

On-stack ABAP Cloud represents a second option when side-by-side extension is not feasible. This pattern uses SAP's cloud-native extension model, which is Clean Core Level A compliant, meaning it is fully supported and upgrade-stable. The trade-offs include more limited flexibility compared to side-by-side approaches and dependency on SAP's extension framework.

Classic extensibility represents a third option, acceptable as a transitional approach but with recognised limitations. Classic extensions may be Level B compliant, meaning they are supported but may require

adjustment during upgrades. Organisations that choose this pattern should document the upgrade risk and plan for eventual migration to more sustainable approaches.

Internal objects represent the least preferred option and should be avoided for AI integration. Extensions built on internal objects carry no compatibility guarantees and are likely to break during upgrades. While such extensions may be technically feasible in the short term, they create technical debt that constrains future flexibility and increases the cost of maintenance.

6.4 A Phased Roadmap for AI-ERP Adoption

The phased roadmap presented here reflects the staged nature of enterprise AI adoption and the practical constraints under which organisations operate. The roadmap is organised into three horizons.

Horizon one, spanning approximately zero to twelve months, focuses on foundation building and low-risk AI adoption. Key activities include data preparation, particularly master data cleansing and governance establishment; identification of low-hanging use cases that are well-bounded and have clear success criteria; foundational governance, including AI principles, risk assessment processes, and accountability structures; and embedding intelligence into high-frequency, rules-based processes where AI can provide immediate value without disrupting core operations.

Horizon two, spanning approximately twelve to twenty-four months, focuses on migration and capability expansion. Key activities include migration to cloud ERP, where feasible, to access the full range of AI capabilities; Clean Core standardisation, including rationalisation of existing extensions; activation of embedded AI scenarios that SAP delivers as standard; and piloting of agentic scenarios in controlled environments to build experience with more autonomous AI capabilities.

Horizon three, spanning approximately twenty-four to thirty-six months and beyond, focuses on scaling and optimisation. Key activities include scaling agentic AI across multiple processes and domains; deployment of industry-specific autonomous scenarios; and establishment of continuous model retraining and monitoring processes. The transition between horizons is not strictly time-bound but depends on the organisation's progress in achieving the prerequisites for each stage.

6.5 Governance Architecture

Governance represents the binding constraint on autonomous enterprise progress. Without adequate governance, organisations cannot safely expand AI autonomy, regardless of technical capability. A governance architecture for AI-ERP integration comprises several elements.

AI governance should be established as a distinct layer above operational and strategic cores, with clear lines of accountability and escalation. This layer should include representation from business functions, IT, legal, compliance, and risk management, ensuring that diverse perspectives inform AI decisions.

Risk registers and transparency artifacts should document AI capabilities, their intended use, their limitations, and the risks they pose. These artifacts support audit, regulatory compliance, and informed

decision-making about AI deployment.

Independent oversight mechanisms should provide challenge and scrutiny of AI decisions, particularly those with significant consequences. This may include internal audit functions, ethics committees, or external reviewers.

Participatory design approaches should involve affected stakeholders in the design and deployment of AI capabilities. This is not merely a matter of consultation but of genuine engagement with those whose work will be affected by AI systems.

Sustainability metrics for computational footprint should be established to monitor and manage the environmental impact of AI operations. These metrics should inform decisions about model selection, deployment architecture, and operational practices.

6.6 Organisational Readiness and Capability Building

Technical frameworks and governance mechanisms are necessary but not sufficient. Organisational readiness and capability building determine whether AI integration succeeds in practice.

Data governance implementation should be prioritised as a foundational activity. This includes establishing clear ownership of data domains, implementing data quality standards and monitoring, and creating processes for data issue resolution. Without effective data governance, AI initiatives will consistently underperform.

Workforce upskilling is essential for the shift from data entry to AI oversight. This includes developing data literacy across the organisation, building skills in prompt engineering and AI interaction, and fostering critical thinking about AI outputs. Training should be ongoing rather than one-time, reflecting the evolving nature of AI capabilities.

Change management for the shift to conversational interfaces requires attention to user experience, support structures, and cultural adaptation. Conversational interfaces change how users interact with ERP systems, and this change may be disorienting for users accustomed to traditional transaction screens. Support structures should include help desks, user communities, and feedback mechanisms.

Cross-functional AI councils can support use case prioritisation and resource allocation. These councils should include representation from business functions, IT, and governance, and should have clear authority to make decisions about AI investment and deployment.

7. Discussion and Synthesis

7.1 Key Findings

This paper has examined the opportunities, challenges, and implementation paths for AI integration in SAP ERP environments. Three findings merit emphasis.

First, the central paradox of AI-ERP integration is that success depends not on technological sophistication but on organisational readiness and architectural discipline. The most sophisticated AI models will fail if deployed on poor-quality data, if users do not trust or adopt them, or if they violate Clean Core principles in ways that destabilise core systems. Conversely, relatively simple AI capabilities can deliver substantial value when deployed on a foundation of data quality and organisational readiness.

Second, the Clean Core tension is real but manageable through appropriate extension patterns. Side-by-side extensibility on SAP BTP and on-stack ABAP Cloud offer viable paths for AI integration that preserve upgrade stability. Classic extensibility and internal object modifications should be avoided or treated as transitional. The key is to recognise that Clean Core is not an impediment to AI value but a prerequisite for sustainable value.

Third, the dual-track strategy reflects the pragmatic reality of mixed ERP landscapes. Organisations cannot wait for full migration to begin realising AI value, but they must be mindful of the constraints that legacy environments impose. Sidecar AI in ECC environments and progressive activation of embedded AI in cloud environments offer a coherent approach to managing this complexity.

7.2 Theoretical Implications

The findings of this paper contribute to theoretical understanding of AI-ERP integration in several ways.

First, the analysis suggests that AI-ERP integration intensifies existing socio-technical challenges rather than creating entirely new ones. The data quality problems, organisational resistance, and governance gaps that have long characterised ERP implementations are amplified when AI is introduced, but they are not fundamentally different in kind. This insight suggests that existing socio-technical theories of ERP implementation remain relevant, though they require extension to address AI-specific dimensions.

Second, the context gap emerges as a critical concept for understanding vendor-customer dynamics in AI-ERP integration. The gap between standardised vendor scenarios and organisation-specific requirements is not merely a technical problem to be solved but a structural feature of the vendor-customer relationship. Recognising this gap helps explain why AI adoption often disappoints and what organisations must do to bridge it.

Third, governance emerges as the binding constraint on autonomous enterprise progress. Technical capability can advance rapidly, but governance frameworks evolve more slowly. Organisations that invest in governance early are better positioned to scale AI autonomy when technical capabilities mature.

7.3 Practical Implications

For practitioners, several implications follow from this analysis.

Clean Core should be understood as a prerequisite for AI value, not an impediment. Organisations that maintain clean cores will be better positioned to adopt AI capabilities as they emerge, while those that

accumulate technical debt will face increasing constraints.

Data readiness should be the highest-priority investment. AI cannot function effectively on poor-quality data, and the investments required to improve data quality are substantial. Organisations should not attempt to deploy AI on data that is not fit for purpose.

The phasing model presented in this paper offers a risk-mitigation strategy. By proceeding through defined horizons, organisations can build capabilities incrementally, learn from experience, and adjust their approaches based on what they learn. Attempting to leap directly to advanced AI autonomy without building foundational capabilities is likely to result in failure.

8. Conclusion and Future Research

8.1 Summary of Contributions

This paper has synthesised fragmented literature on AI-ERP integration into a coherent framework addressing lifecycle and socio-technical dimensions. It has identified the Clean Core tension and the context gap as structural challenges distinguishing AI-ERP integration from traditional ERP customisation. It has proposed a dual-track implementation strategy, a DevOps framework, Clean Core-compatible extension patterns, a phased roadmap, and a governance architecture.

The paper's central argument is that successful AI integration depends not on technological sophistication but on establishing a sustainable balance between the stability requirements of ERP and the adaptive demands of AI. This balance is achieved through architectural discipline, data readiness, governance mechanisms, and organisational capability building.

8.2 Limitations and Boundaries of the Framework

The framework developed here requires empirical validation. The reliance on published literature and vendor documentation, rather than primary data, means the framework's applicability across different organisational contexts remains to be tested. The rapid pace of change in AI and ERP technologies creates obsolescence risk; specific product names and capabilities may evolve even as underlying principles remain stable.

8.3 Future Research Directions

Several directions for future research follow from this analysis.

Empirical validation of the phased roadmap through longitudinal case studies would test the framework's applicability and refine its components. Such studies could examine how organisations progress through the horizons, what challenges they encounter, and what factors influence success.

Development of AI-ERP-specific maturity models would provide organisations with tools for assessing their current state and planning their development. Such models could draw on established maturity

frameworks from related domains while addressing AI-specific dimensions.

Comparative analysis across ERP vendors would examine whether the challenges and opportunities identified here are specific to SAP or generalisable to other platforms. This analysis could inform vendor selection decisions and identify best practices across the industry.

Investigation of regional variation in governance requirements would examine how different regulatory frameworks shape AI-ERP integration. Such research could inform the development of governance frameworks that are both effective and adaptable to different jurisdictional contexts.

Longitudinal studies of workforce transformation in autonomous enterprises would examine how roles, skills, and organisational structures evolve as AI capabilities mature. Such research could inform workforce planning and change management practices.

References

1. Almeida F, Santos J, Monteiro J (2020) The challenges and opportunities in the digitalization of companies in a post-COVID-19 world. IEEE Eng Manage Rev 48(3):97–103
2. Chenna K (2024) Revolutionizing Supply Chain: The Synergy of Warehouse Management and Data Analytics. Int J Sci Res Comput Sci Eng Inform Technol 10(6):122–131
3. Chenna K (2024) The Convergence of SAP PP, RPA, and IBP: A Comprehensive analysis of integration Benefits and Challenges in Modern manufacturing. Int J Multidisciplinary Res, *6*(6)
4. Katuu S (2020) Enterprise resource planning: Past, present, and future. New Rev Inform Netw 25(1):37–46
5. Russell S, Norvig P (2021) *Artificial intelligence: A modern approach* (4th ed.). Pearson
6. SAP SE (2022) Clean core: Extensibility guide for SAP S/4HANA. SAP
7. SAP SE (2023) Business AI: Transforming enterprise applications. SAP
8. Webster J, Watson RT (2002) Analyzing the past to prepare for the future: Writing a literature review. MIS Q 26(2):xiii–xxiii