British Journal of Multidisciplinary Studies
ISSN: 2978-8110
DOI: 10.32996/bjmss
Journal Homepage: al-kindipublisher.com/index.php/bjmss/index

BJMSS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

RESEARCH ARTICLE

Enterprise Governance for AI-Driven SAP Landscapes: Security, Compliance, and Responsible AI

Purnachander Rao, Nadipelli
IT Technical Manager | ERP Transformation | SAP S/4HANA, BTP - Corning Inc.
Corresponding Author: Purnachander Rao, Nadipelli, **E-mail:** yuvachandramarasani@gmail.com

ABSTRACT

The rapid integration of artificial intelligence into SAP-enabled enterprise environments is transforming business processes, decision-making, automation, and data-driven operations. However, the increasing use of AI within complex SAP landscapes also introduces significant governance challenges involving cybersecurity, regulatory compliance, data privacy, algorithmic transparency, accountability, and ethical responsibility. This study examines enterprise governance mechanisms required to support secure, compliant, and responsible AI adoption across SAP ecosystems. It integrates established perspectives on AI governance, explainable artificial intelligence, role-based access control, data governance, continuous auditing, and responsible AI principles to develop a structured governance approach for enterprise environments. Particular attention is given to identity and access management, authorization controls, segregation of duties, explainability, auditability, privacy protection, bias mitigation, and human oversight. The study further proposes an integrated governance framework that aligns technological controls with organizational responsibilities, regulatory expectations, and ethical safeguards. Practical implementation guidelines are presented to help enterprises balance AI-enabled innovation with operational resilience, accountability, and business trust. The findings emphasize that effective AI governance must be continuous, risk-based, transparent, and embedded within the broader SAP governance architecture.

KEYWORDS

Artificial intelligence governance, SAP, enterprise governance, responsible AI, cybersecurity, regulatory compliance, explainable AI, data privacy, auditability, access control.

ARTICLE INFORMATION

ACCEPTED: 10 February 2024 **PUBLISHED:** 20 March 2024 **DOI:** 10.32996/bjmss.2024.3.2.4

1. Introduction

Artificial intelligence has become an increasingly important component of contemporary enterprise systems, reshaping how organizations automate operations, analyze information, manage resources, and support strategic decision-making. Within SAP environments, AI capabilities are being embedded across enterprise resource planning, supply chain management, finance, human capital management, customer relationship management, procurement, and business analytics. This transition reflects the broader evolution of enterprise systems from transactional platforms toward intelligent, data-driven ecosystems capable of prediction, recommendation, automation, and adaptive decision support. As organizations increasingly integrate machine learning, intelligent automation, predictive analytics, and generative AI into SAP-enabled processes, enterprise governance must evolve accordingly. Traditional governance models that primarily focused on system configuration, access control, segregation of duties, and regulatory compliance are no longer sufficient to address the risks introduced by autonomous or semi-autonomous AI-driven processes.

AI integration creates a multidimensional governance challenge because risks can emerge from data, algorithms, system access, decision logic, organizational processes, and regulatory obligations simultaneously. AI-enabled SAP environments may process highly sensitive financial, employee, supplier, operational, and customer information, making data privacy, identity management, authorization, and cybersecurity critical governance priorities. At the same time, AI systems can generate recommendations or decisions whose reasoning may be difficult to interpret, creating concerns regarding explainability, accountability, bias, and

AL-KINDI

Copyright: © 2024 the Author(s). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC-BY) 4.0 license (https://creativecommons.org/licenses/by/4.0/). Published by Al-Kindi Centre for Research and Development, London, United Kingdom.

auditability. Research on trustworthy and responsible AI consistently emphasizes transparency, fairness, accountability, human oversight, and effective risk management as central requirements for sustainable AI adoption (Thiebes et al., 2021; Mäntymäki et al., 2022; Díaz-Rodríguez et al., 2023). Explainability is particularly important in enterprise contexts because organizations must often demonstrate why automated outcomes occurred and whether those outcomes comply with established policies, regulations, and internal controls (Ali et al., 2023).

The central research problem addressed in this study is the absence of a sufficiently integrated governance approach that connects established enterprise controls with emerging responsible AI requirements within complex SAP landscapes. Existing studies provide valuable guidance on AI governance, ethical AI, data governance, algorithmic auditing, access control, and cloud enterprise resource planning security, yet these perspectives are often treated separately. This fragmentation creates practical difficulties for organizations attempting to govern AI systems embedded within interconnected enterprise applications. The objective of this study is therefore to develop a comprehensive governance perspective that integrates security, identity and authorization management, regulatory compliance, data privacy, explainability, auditability, ethical principles, and human accountability into a coherent framework for AI-driven SAP environments.

The significance of the study lies in its emphasis on governance as an organizational and technical capability rather than a purely compliance-oriented activity. Effective governance should enable innovation while maintaining business trust, operational resilience, and responsible decision-making. As Birkstedt et al. (2023) and Papagiannidis et al. (2023) indicate, AI governance increasingly requires formal structures, clearly defined responsibilities, continuous monitoring, and organizational mechanisms capable of translating ethical principles into operational practice. The scope of this study therefore focuses on enterprise-level governance of AI capabilities operating within or alongside SAP landscapes. It examines security, authorization, privacy, compliance, explainability, auditability, responsible AI, and implementation practices, while emphasizing practical governance mechanisms that organizations can incorporate into existing SAP control architectures.

2. AI-Driven SAP Landscapes and Enterprise Governance

The evolution of SAP from a conventional enterprise resource planning platform toward an intelligent enterprise ecosystem reflects a broader transformation in how organizations use digital technologies to coordinate business processes, manage data, and support decision-making. Traditional SAP implementations were largely designed around structured transaction processing, standardized workflows, financial control, supply chain coordination, human resource administration, procurement, and reporting. Over time, however, enterprise environments have increasingly incorporated cloud computing, advanced analytics, machine learning, robotic process automation, predictive modelling, and intelligent decision-support capabilities. This shift has changed the role of SAP from being primarily a system of record into a broader digital environment in which data can be analysed continuously and used to automate or augment operational and managerial decisions. The emergence of AI-enabled enterprise environments therefore represents not simply an additional technological layer, but a fundamental change in how organizational authority, data, processes, and decisions interact.

In an AI-driven SAP landscape, enterprise applications may combine transactional data with analytical models capable of identifying patterns, predicting outcomes, detecting anomalies, recommending actions, and automating repetitive activities. Machine learning can support areas such as demand forecasting, fraud detection, predictive maintenance, cash-flow forecasting, customer segmentation, procurement optimization, and workforce analytics. Automation can further connect AI-generated insights with operational workflows, enabling organizations to respond to information at greater speed and scale. These developments can produce significant efficiency and decision-quality benefits, yet they also make enterprise systems more complex. Decisions that were once clearly attributable to specific employees or predefined business rules may increasingly depend on statistical models, automated recommendations, and interconnected data pipelines. Consequently, organizations must govern not only who can access an SAP system, but also what data AI systems can use, how models reach conclusions, who can approve automated decisions, and how those decisions can later be reconstructed and audited.

The increasing importance of explainability illustrates this governance challenge. AI models can generate accurate outputs while providing limited visibility into how individual predictions or recommendations were produced. In enterprise environments, such opacity may become problematic when AI influences financially significant, legally sensitive, or operationally critical decisions. Explainable artificial intelligence therefore seeks to provide understandable accounts of model behaviour and specific outputs so that users, auditors, regulators, and decision-makers can evaluate whether an AI-supported outcome is reliable and appropriate (Arrieta et al., 2020; Vilone & Longo, 2021). Recent scholarship further emphasizes that explainability is not merely a technical property but an essential component of trustworthy AI because it supports transparency, accountability, contestability, and informed human oversight (Ali et al., 2023). Within an SAP context, these requirements become especially important when AI outputs interact directly with business transactions, approvals, master data, financial controls, or other enterprise workflows.

Enterprise AI governance can therefore be understood as the system of structures, policies, responsibilities, controls, and monitoring mechanisms used to direct and oversee organizational AI. Mäntymäki et al. (2022) characterize organizational AI

governance as the mechanisms through which organizations align AI development and use with strategic objectives, ethical expectations, and accountability requirements. Birkstedt et al. (2023) similarly demonstrate that AI governance extends beyond technology management and requires coordination among organizational, regulatory, ethical, and technical domains. In SAP landscapes, this means AI governance should not function independently of established enterprise governance. Instead, it should be integrated with existing identity management, authorization models, segregation-of-duties controls, data governance, cybersecurity, risk management, internal auditing, privacy management, and compliance structures.

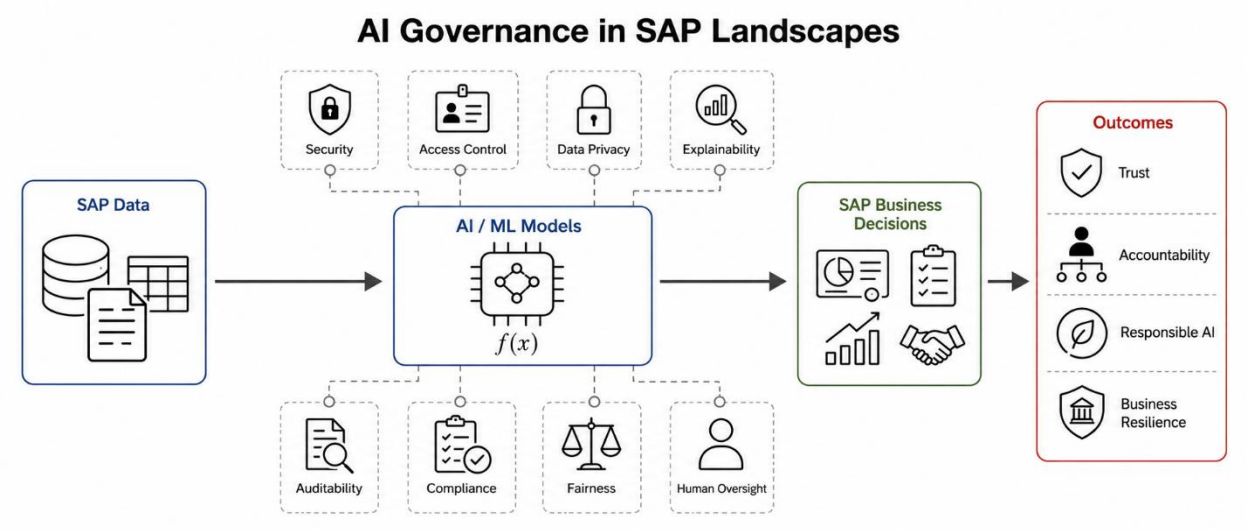
This integration is important because AI-enabled SAP environments introduce several interconnected categories of governance risk. One major risk concerns data. AI performance depends heavily on the quality, representativeness, legality, and integrity of the data used for training, validation, and operational inference. Poor-quality or biased enterprise data may produce inaccurate or discriminatory outputs, while inappropriate data access can create privacy and confidentiality violations. Janssen et al. (2020) argue that effective data governance is fundamental to trustworthy AI because organizations must clearly establish data ownership, accountability, provenance, quality controls, and appropriate conditions of use. Dataset documentation mechanisms can further improve transparency by recording how data were created, what populations they represent, and what limitations may affect their use (Gebru et al., 2021).

A second governance risk involves identity, authorization, and access privileges. SAP environments typically contain sensitive financial, employee, supplier, customer, and operational information. When AI systems are connected to these datasets and business processes, poorly governed access can amplify the consequences of unauthorized activity. Traditional role-based access control remains highly relevant because access rights should correspond to legitimate organizational responsibilities rather than unrestricted technical capability (Sandhu et al., 1996; Ferraiolo et al., 2001). However, AI introduces additional considerations, including access to model outputs, training datasets, prompts, automated actions, and model configuration. Governance should therefore determine not only who can access enterprise data but who may deploy, modify, approve, override, or act upon AI systems.

A third risk concerns accountability and auditability. When AI contributes to enterprise decisions, organizations must retain sufficient evidence to determine what occurred, which data were used, what model or rule generated the output, who approved the action, and whether appropriate controls were followed. Algorithmic auditing provides one mechanism for reducing this accountability gap by evaluating AI systems throughout their lifecycle rather than only after deployment (Raji et al., 2020). This principle complements continuous auditing practices traditionally associated with ERP environments, where automated controls and ongoing monitoring can improve the identification of irregularities and control failures (Kuhn & Sutton, 2010).

Finally, AI-driven SAP landscapes raise ethical and organizational risks involving bias, fairness, transparency, responsibility, and human oversight. Bias can originate from historical data, feature selection, model design, or organizational practices and may lead to systematically unequal outcomes (Mehrabi et al., 2021). Ethical principles alone are insufficient unless organizations translate them into operational controls, measurable responsibilities, and enforceable procedures (Mittelstadt, 2019; Morley et al., 2020). Effective enterprise governance must therefore combine technical safeguards with human accountability. The objective is not to prevent AI-enabled innovation, but to ensure that innovation occurs within a controlled environment where security, compliance, transparency, fairness, and business responsibility are treated as integral components of enterprise architecture. In this sense, AI governance becomes a strategic capability that enables organizations to exploit intelligent SAP technologies while preserving organizational trust and control.

Figure 1. AI Governance in SAP Landscapes



Conceptual framework showing how SAP data flows through AI/ML models under key governance controls, including security, access control, data privacy, explainability, auditability, compliance, fairness, and human oversight, to support trusted, accountable, responsible, and resilient SAP business decisions.

3. Security, Identity, and Authorization Governance

Security governance is a foundational requirement for AI-driven SAP landscapes because artificial intelligence expands both the capabilities and the potential attack surface of enterprise systems. SAP environments frequently contain highly sensitive financial, employee, procurement, customer, operational, and strategic data. Once AI models, intelligent automation, and analytics are connected to these resources, security governance must extend beyond conventional authentication and system access. Organizations must also govern who can access AI models, training data, prompts, outputs, automated workflows, configuration settings, and decision interfaces. Effective security governance therefore requires a coordinated framework combining identity management, authorization, privileged-access controls, segregation of duties, cybersecurity monitoring, and continuous risk assessment.

Identity and Access Management (IAM) provides the first layer of protection by ensuring that every user, application, service, and AI component interacting with the SAP ecosystem has a verified identity and an appropriate level of access. IAM should incorporate strong authentication, user provisioning, access certification, lifecycle management, and timely removal of obsolete privileges. Access should be governed according to the principle of least privilege, meaning that users and systems receive only the permissions necessary to perform approved tasks. This is particularly important in AI-enabled environments because excessive access can allow automated systems to retrieve sensitive information or initiate actions beyond their intended purpose. Database security research has long emphasized the importance of controlled authentication, authorization, and access enforcement in protecting enterprise information assets (Bertino & Sandhu, 2005).

Role-Based Access Control (RBAC) complements IAM by assigning permissions according to organizational roles rather than granting rights independently to each user. RBAC is especially suitable for complex SAP environments because permissions can be aligned with business responsibilities such as finance officer, procurement manager, system administrator, auditor, or data analyst. The fundamental principle is that access should correspond to defined job responsibilities and organizational authority (Sandhu et al., 1996; Ferraiolo et al., 2001). In AI-enabled SAP landscapes, RBAC should be extended to AI-specific activities. For example, authorization should distinguish between users who can view AI recommendations, modify model parameters, access training datasets, approve automated decisions, or override AI-generated outputs.

Privileged access represents a particularly sensitive governance area. Administrators, developers, security specialists, and certain technical service accounts may possess extensive permissions capable of modifying business rules, system configurations, datasets, or AI workflows. Such privileges should therefore be tightly controlled, monitored, and periodically reviewed. Privileged Access Management should include approval workflows, session monitoring, time-limited access, credential protection, and detailed activity logging. Segregation of Duties (SoD) is equally important because critical enterprise functions should not be concentrated under a single identity. For example, an individual who configures an AI-enabled payment process should not also

possess unrestricted authority to approve and execute payments. Such separation reduces the risk of fraud, manipulation, unauthorized model changes, and control circumvention.

AI integration also creates new cybersecurity threats. Attackers may target models, data pipelines, application programming interfaces, automation interfaces, or identity systems. Threats can include data poisoning, adversarial manipulation, model theft, unauthorized inference, prompt-based exploitation, privilege escalation, and deliberate modification of training or operational data. AI systems may also unintentionally expose confidential enterprise information through poorly controlled outputs. These threats increase the importance of secure system design, encryption, access logging, vulnerability management, secure interfaces, data validation, and continuous testing. Salih et al. (2021) emphasize that security remains one of the critical issues surrounding cloud ERP adoption, especially as enterprise environments become more interconnected and dependent on external services.

Continuous security monitoring is therefore essential. Governance cannot rely solely on controls established during initial system deployment because user privileges, AI models, threats, data flows, and business processes continually change. Organizations should monitor authentication events, failed access attempts, privilege changes, unusual transactions, AI-model activities, data movements, configuration changes, and exceptions to established controls. Continuous auditing approaches in ERP systems provide a useful foundation by allowing organizations to detect anomalies and control failures more rapidly than periodic manual reviews (Kuhn & Sutton, 2010). AI-enabled monitoring can further improve detection by identifying patterns that may indicate unauthorized or abnormal behaviour.

Risk management should operate throughout the entire AI lifecycle, from system design and data acquisition to deployment, monitoring, modification, and retirement. Security risks should be assessed according to likelihood, potential impact, affected business processes, regulatory consequences, and availability of mitigating controls. Responsible AI governance therefore requires security controls to be integrated with organizational accountability rather than treated as a purely technical responsibility. Effective security, identity, and authorization governance ultimately ensures that AI capabilities enhance SAP operations without undermining confidentiality, integrity, availability, regulatory compliance, or organizational trust.

Table 1. Core Security and Authorization Controls in AI-Driven SAP Landscapes

Governance Area	Primary Control	Main Risk Addressed	Expected Outcome
Identity Management	Strong authentication and identity verification	Unauthorized access	Verified users and services
Access Management	Least-privilege access	Excessive permissions	Controlled system access
RBAC	Role-based permissions	Inappropriate access rights	Access aligned with job roles
Privileged Access	Approval, logging, session control	Administrative abuse	Controlled high-risk access
Segregation of Duties	Separation of critical functions	Fraud and conflicts of interest	Strong internal control
AI Security	Model and data protection	Manipulation and leakage	Secure AI operations
Monitoring	Continuous logging and anomaly detection	Undetected threats	Early threat identification
Risk Management	Periodic and continuous assessment	Emerging vulnerabilities	Adaptive security governance

Figure 2. Security Governance Structure for AI-Driven SAP Landscapes

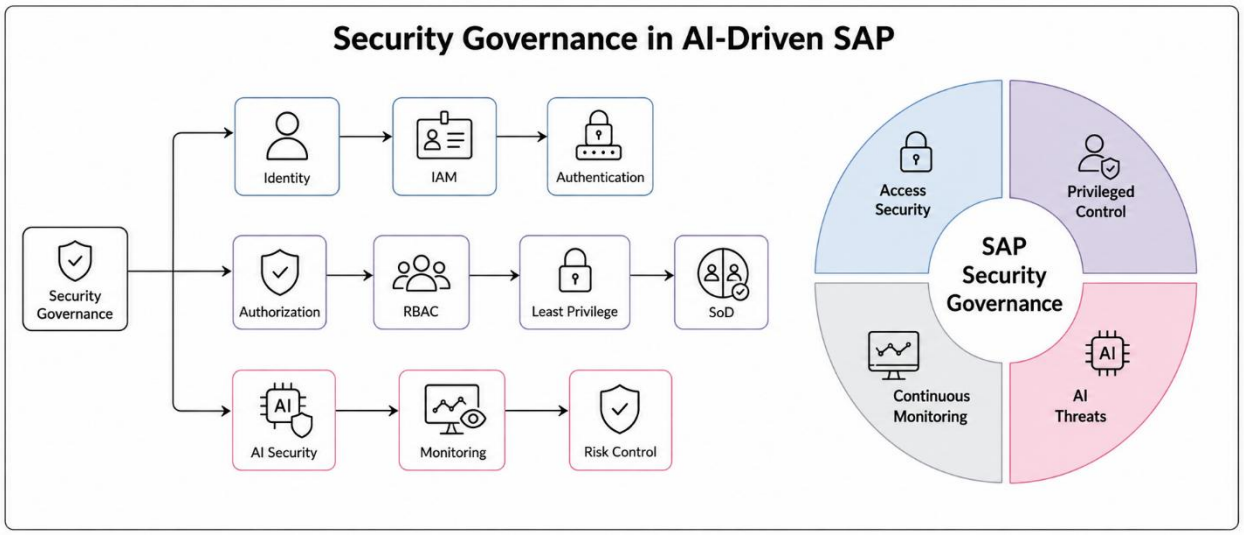


Figure 2 illustrates the integration of identity management, role-based authorization, privileged access control, segregation of duties, AI security, continuous monitoring, and risk management within an AI-enabled SAP environment. The framework shows how these controls collectively support secure access, threat detection, accountability, and continuous governance across enterprise systems.

4. Regulatory Compliance, Data Privacy, and Auditability

The integration of artificial intelligence into SAP landscapes creates significant regulatory and data-governance responsibilities because enterprise systems routinely process large volumes of financial, employee, customer, supplier, and operational information. AI extends the value of these datasets by enabling prediction, automation, classification, and decision support, but it also increases the complexity of determining how information is collected, processed, shared, retained, and used. Regulatory compliance in an AI-driven SAP environment must therefore be embedded throughout the AI and data lifecycle rather than treated as an isolated legal or technical activity. Governance mechanisms must connect privacy protection, data quality, accountability, documentation, explainability, and continuous auditing to ensure that AI-supported activities remain consistent with organizational policies and applicable regulatory obligations.

The General Data Protection Regulation (GDPR) provides an important foundation for enterprise data protection. Its principles of lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity, confidentiality, and accountability have direct implications for AI-enabled enterprise systems. Organizations operating SAP environments must understand what personal information is processed, why it is required, where it is stored, how long it is retained, and which users or AI applications have access to it. Tikkinen-Piri et al. (2018) emphasize that GDPR substantially strengthened organizational responsibilities concerning personal-data processing and accountability. Within AI-enabled SAP systems, privacy governance must consequently extend to datasets used for model development, training, testing, inference, and automated decision processes. Sensitive information should not become available to an AI application merely because it exists elsewhere within the enterprise environment.

Effective data governance is therefore fundamental to regulatory compliance. Data should be governed throughout its lifecycle, encompassing collection, classification, storage, processing, sharing, archival, and deletion. Ownership and stewardship responsibilities should be explicitly assigned, while data lineage should enable organizations to establish where information originated and how it has been transformed. Janssen et al. (2020) argue that trustworthy AI depends substantially on effective data governance because poor-quality, inadequately controlled, or poorly understood datasets undermine the reliability and legitimacy of AI outcomes. Dataset documentation can further strengthen governance by recording intended uses, composition, limitations, collection procedures, and potential biases (Gebru et al., 2021). For SAP environments, these practices provide an essential connection between enterprise master data, transactional information, analytical datasets, and AI models.

Auditability represents another central requirement. Organizations must be capable of reconstructing significant AI-supported actions after they occur. An effective AI audit trail should record the relevant user or system identity, data source, model version, input, output, timestamp, authorization status, human intervention, and resulting business action. This information allows auditors

and governance teams to establish why a decision occurred and who ultimately held responsibility. Raji et al. (2020) demonstrate the importance of end-to-end algorithmic auditing as a mechanism for closing accountability gaps throughout the AI lifecycle. This approach is particularly relevant to SAP, where continuous auditing has historically provided mechanisms for identifying anomalous transactions and control weaknesses within highly integrated ERP environments (Kuhn & Sutton, 2010).

Regulatory compliance should consequently be translated into enforceable SAP controls rather than remaining primarily documentation-based. Policies should be mapped to authorization rules, approval workflows, retention requirements, data classifications, monitoring mechanisms, and audit evidence. Automated compliance checks can identify unauthorized access, unusual data transfers, expired privileges, missing approvals, or deviations from established policies. Nevertheless, automation does not eliminate human accountability. High-impact exceptions and sensitive AI-supported decisions should remain subject to appropriate review and escalation.

Explainability and transparency complete this governance structure. Organizations must be able to provide meaningful information about how important AI-supported outcomes were generated, particularly where such outcomes affect employees, customers, financial transactions, or compliance obligations. Explainability mechanisms should therefore be accompanied by model documentation, data records, validation evidence, performance metrics, limitations, change histories, and clearly assigned ownership. Ali et al. (2023) identify explainability as a central component of trustworthy AI, while Díaz-Rodríguez et al. (2023) connect transparency, accountability, human oversight, and regulatory requirements within responsible AI systems. In AI-driven SAP landscapes, regulatory compliance, data privacy, explainability, and auditability should therefore operate as interconnected governance capabilities that create a traceable chain from data acquisition to AI decision and ultimately to accountable business action.

Table 2. Key Compliance and Data Governance Controls

Governance Area	Key Control	Purpose
Data Privacy	Minimization and access control	Protect personal data
Data Lifecycle	Classification and retention	Control data use
Auditability	Logs and model records	Ensure traceability
Compliance	Policy-to-control mapping	Meet obligations
Explainability	Model and decision documentation	Support transparency
Accountability	Defined ownership and review	Assign responsibility

Figure 3. Theoretical Relationship Between AI Governance Maturity and Compliance Risk

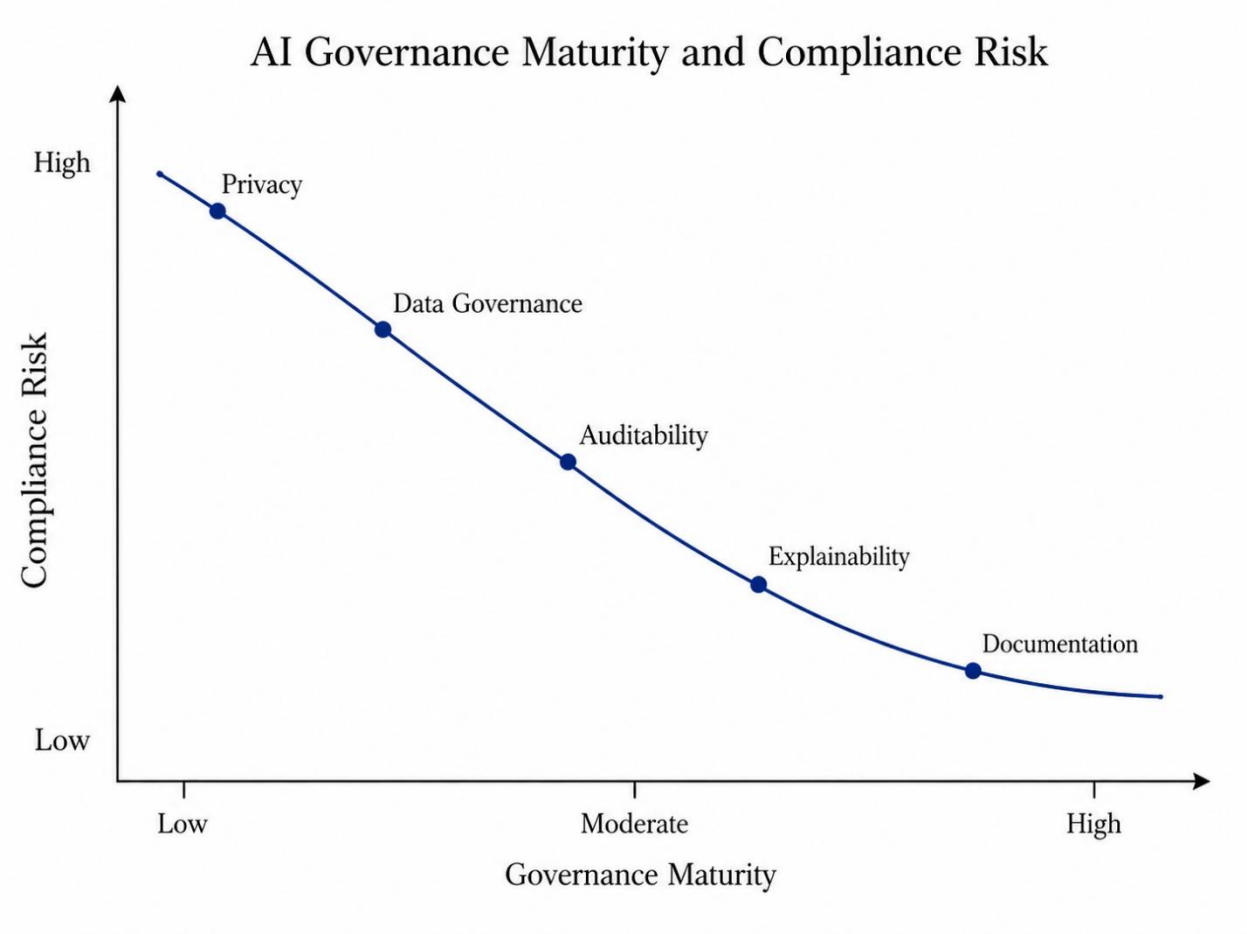


Figure 3 illustrates the expected inverse relationship between governance maturity and compliance risk in AI-driven SAP environments. As organizations strengthen data privacy, data governance, auditability, explainability, documentation, and accountability mechanisms, the likelihood of regulatory violations, uncontrolled data use, opaque AI decisions, and governance failures is expected to decline.

5. Responsible AI Governance Framework for SAP

Responsible artificial intelligence governance within SAP environments requires organizations to move beyond conventional technology controls and establish mechanisms that ensure AI systems remain fair, transparent, accountable, secure, and aligned with organizational values. Because AI-enabled SAP applications may influence financial decisions, procurement, workforce management, customer services, supply-chain operations, and other critical processes, governance must address both the technical performance of AI and the consequences of its outputs. Responsible AI should therefore be embedded throughout the complete system lifecycle, from data preparation and model development to deployment, monitoring, modification, and retirement.

Fairness and bias mitigation represent fundamental components of this framework. Bias may originate from historical enterprise data, inappropriate variables, unrepresentative datasets, or model-development decisions. Such biases can produce systematically disadvantageous outcomes, particularly when AI supports recruitment, employee evaluation, credit-related assessment, supplier selection, or customer segmentation. Organizations should consequently conduct periodic bias assessments, examine model performance across relevant groups, validate training datasets, and document corrective measures. Mehrabi et al. (2021) demonstrate that bias can emerge from multiple stages of the machine-learning lifecycle, reinforcing the need for continuous rather than one-time fairness assessment.

Transparency and explainability are equally important. AI-generated recommendations should be sufficiently interpretable for relevant users to understand the principal factors influencing significant outputs. Explainable AI mechanisms can strengthen trust

while enabling managers, auditors, and compliance personnel to examine whether decisions reflect acceptable business logic (Arrieta et al., 2020; Ali et al., 2023). Within SAP systems, explanations should be supported by model documentation, input records, decision logs, performance information, and defined limitations. Transparency should therefore apply not only to model outputs but also to data provenance, model ownership, intended use, and decision authority.

Human oversight provides the bridge between automated intelligence and organizational accountability. AI systems should support rather than eliminate responsibility for high-impact decisions. Critical AI-generated actions should contain clearly defined approval, review, escalation, and override mechanisms. Human reviewers must have sufficient authority and information to challenge inappropriate outcomes. This reflects broader trustworthy AI principles emphasizing human agency, accountability, fairness, transparency, and harm prevention (Floridi et al., 2018; Thiebes et al., 2021; Díaz-Rodríguez et al., 2023).

Based on these principles, this study proposes an integrated SAP AI governance framework organized around six interconnected pillars: data governance, security and authorization, fairness, explainability, human oversight, and continuous assurance. These controls should surround the complete AI lifecycle rather than operate as separate compliance activities. Data owners should ensure data quality and appropriate use; AI developers should document models and conduct technical validation; business process owners should define acceptable use and performance expectations; security teams should govern access and cyber risk; compliance and legal functions should interpret regulatory obligations; internal auditors should independently assess controls; and senior governance committees should establish risk appetite and accountability.

The proposed framework therefore treats responsible AI governance as a shared enterprise responsibility. Effective SAP AI governance depends on clearly assigned ownership, documented controls, measurable standards, continuous monitoring, and formal escalation mechanisms. Such integration enables organizations to pursue AI-driven innovation while preserving fairness, transparency, regulatory accountability, and organizational trust.

6. Implementation Guidelines

Effective implementation of AI governance within SAP environments requires a structured and phased approach that aligns technology, business processes, regulatory obligations, and organizational accountability. The first implementation priority should be to establish a formal AI governance structure with clearly defined ownership. Organizations should create a multidisciplinary governance committee involving SAP administrators, cybersecurity specialists, data owners, compliance officers, legal advisers, internal auditors, AI developers, and business process owners. This committee should define governance policies, risk tolerance, approval requirements, and escalation procedures for AI-enabled SAP applications. Clear assignment of responsibilities is essential because AI governance failures often result from fragmented ownership rather than from technical deficiencies alone (Mäntymäki et al., 2022; Birkstedt et al., 2023).

The second step is to develop an enterprise inventory of AI-enabled SAP use cases. Each AI system should be documented according to its business purpose, data sources, model type, users, automated actions, affected stakeholders, and potential consequences. Organizations should classify applications according to risk, with higher-risk systems receiving stronger validation, approval, monitoring, and human oversight. This risk-based approach enables governance resources to be concentrated where potential financial, operational, ethical, or regulatory harm is greatest.

Data and access controls should then be integrated directly into SAP architecture. Organizations should enforce role-based access, least privilege, segregation of duties, privileged-access monitoring, and periodic access certification. AI training and operational datasets should undergo quality assessment, privacy review, lineage tracking, and retention control. Sensitive data should be accessed only for clearly defined purposes and protected through appropriate technical and organizational safeguards.

Operational deployment should also include model validation, explainability assessment, bias testing, cybersecurity evaluation, and documented human-review procedures. Model changes should follow formal change-management processes, with version histories and approval records maintained for audit purposes. Continuous monitoring should measure model performance, unusual access patterns, policy violations, bias indicators, and compliance exceptions. Raji et al. (2020) support lifecycle-based auditing as a mechanism for strengthening algorithmic accountability.

Finally, organizations should conduct periodic governance reviews and employee training. Governance policies must evolve as AI systems, threats, regulations, and business requirements change. Successful implementation therefore depends on treating AI governance as a continuous enterprise capability rather than a one-time compliance exercise.

7. Discussion

The analysis demonstrates that AI governance within SAP landscapes cannot be effectively achieved through isolated security, compliance, or ethical controls. The central requirement is integration. Identity management, authorization, data governance, privacy, explainability, fairness, auditing, and human oversight must operate as complementary components of a unified enterprise governance architecture. This finding is consistent with the broader literature, which increasingly presents trustworthy AI as a multidimensional organizational challenge rather than a purely technical problem (Thiebes et al., 2021; Birkstedt et al., 2023).

A major implication is that existing SAP governance mechanisms provide a valuable foundation but require extension to address AI-specific risks. Traditional controls such as role-based access, segregation of duties, audit logging, and data governance remain important; however, they must now account for model behaviour, automated recommendations, training data, algorithmic bias, model changes, and explainability. Organizations that merely add AI capabilities to existing SAP systems without adapting governance structures may increase operational efficiency while simultaneously creating new accountability and compliance weaknesses.

The proposed framework also highlights an important balance between innovation and control. Excessive governance may slow experimentation, while insufficient governance may expose organizations to security, ethical, regulatory, and reputational risks. A risk-based governance model provides a practical middle ground by applying stronger oversight to high-impact AI applications while permitting proportionate controls for lower-risk use cases. Governance should therefore function as an enabler of sustainable AI adoption rather than as a barrier to technological innovation.

8. Conclusion

AI-driven SAP landscapes offer organizations significant opportunities for automation, predictive decision-making, operational efficiency, and intelligent enterprise management. However, these benefits introduce new risks involving cybersecurity, data privacy, access control, algorithmic bias, explainability, regulatory compliance, and accountability. This study argues that effective governance must integrate traditional SAP controls with responsible AI principles across the complete AI lifecycle.

A comprehensive governance approach should combine identity and authorization management, segregation of duties, data governance, continuous auditing, explainability, fairness assessment, human oversight, and clearly defined organizational responsibility. Governance should also remain continuous and risk-based, allowing controls to evolve alongside changes in AI models, enterprise processes, regulatory expectations, and emerging security threats.

Ultimately, responsible AI governance is essential for transforming AI from an experimental capability into a trusted enterprise asset. Organizations that embed governance into SAP architecture can accelerate AI adoption while maintaining security, accountability, compliance, operational resilience, and long-term stakeholder trust.

Funding: Funding information was not supplied in the manuscript.

Conflicts of Interest: A conflict-of-interest declaration was not supplied in the manuscript.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J. M., Confalonieri, R., Guidotti, R., Del Ser, J., Díaz-Rodríguez, N., & Herrera, F. (2023). Explainable artificial intelligence (XAI): What we know and what is left to attain trustworthy artificial intelligence. Information Fusion, 99, 101805. DOI: 10.1016/j.inffus.2023.101805

Angelov, P. P., Soares, E. A., Jiang, R., Arnold, N. I., & Atkinson, P. M. (2021). Explainable artificial intelligence: An analytical review. WIREs Data Mining and Knowledge Discovery, 11(5), e1424. DOI: 10.1002/widm.1424

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. Information Fusion, 58, 82–115. DOI: 10.1016/j.inffus.2019.12.012

Bertino, E., & Sandhu, R. (2005). Database security—Concepts, approaches, and challenges. IEEE Transactions on Dependable and Secure Computing, 2(1), 2–19. DOI: 10.1109/TDSC.2005.9

Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. (2023). AI governance: Themes, knowledge gaps and future agendas. Internet Research, 33(7), 133–167. DOI: 10.1108/INTR-01-2022-0042

Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López de Prado, M., Herrera-Viedma, E., & Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. Information Fusion, 99, 101896. DOI: 10.1016/j.inffus.2023.101896

Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., & Chandramouli, R. (2001). Proposed NIST standard for role-based access control. ACM Transactions on Information and System Security, 4(3), 224–274. DOI: 10.1145/501978.501980

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and Machines, 28, 689–707. DOI: 10.1007/s11023-018-9482-5

Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. Frontiers in Artificial Intelligence, 5, 779799. DOI: 10.3389/frai.2022.779799

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. Communications of the ACM, 64(12), 86–92. DOI: 10.1145/3458723

Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. Minds and Machines, 30, 99–120. DOI: 10.1007/s11023-020-09517-8

Janssen, M., Brous, P., Estevez, E., Barbosa, L. S., & Janowski, T. (2020). Data governance: Organizing data for trustworthy artificial intelligence. Government Information Quarterly, 37(3), 101493. DOI: 10.1016/j.giq.2020.101493

Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. Nature Machine Intelligence, 1, 389–399. DOI: 10.1038/s42256-019-0088-2

Kuhn, J. R., Jr., & Sutton, S. G. (2010). Continuous auditing in ERP system environments: The current state and future directions. Journal of Information Systems, 24(1), 91–112. DOI: 10.2308/jis.2010.24.1.91

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys, 54(6), Article 115, 1–35. DOI: 10.1145/3457607

Minkkinen, M., & Mäntymäki, M. (2023). Discerning between the "easy" and "hard" problems of AI governance. IEEE Transactions on Technology and Society, 4(2), 188–194.

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. Proceedings of the Conference on Fairness, Accountability, and Transparency, 220–229. DOI: 10.1145/3287560.3287596

Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. Nature Machine Intelligence, 1, 501–507. DOI: 10.1038/s42256-019-0114-4

Morley, J., Floridi, L., Kinsey, L., & Elhalal, A. (2020). From what to how: An initial review of publicly available AI ethics tools, methods and research to translate principles into practices. Science and Engineering Ethics, 26, 2141–2168. DOI: 10.1007/s11948-019-00165-5

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. AI and Ethics, 2, 603–609. DOI: 10.1007/s43681-022-00143-x

Papagiannidis, E., Enholm, I. M., Dremel, C., Mikalef, P., & Krogstie, J. (2023). Toward AI governance: Identifying best practices and potential barriers and outcomes. Information Systems Frontiers, 25, 123–141. DOI: 10.1007/s10796-022-10251-y

Radclyffe, C., Ribeiro, M., & Wortham, R. H. (2023). The assessment list for trustworthy artificial intelligence: A review and recommendations. Frontiers in Artificial Intelligence, 6, 1020592. DOI: 10.3389/frai.2023.1020592

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 33–44. DOI: 10.1145/3351095.3372873

Salih, S., Hamdan, M., Abdelmaboud, A., Abdelaziz, A., Abdelsalam, S., Althobaiti, M. M., Cheikhrouhou, O., Hamam, H., & Alotaibi, F. (2021). Prioritising organisational factors impacting cloud ERP adoption and the critical issues related to security, usability, and vendors: A systematic literature review. Sensors, 21(24), 8391. DOI: 10.3390/s21248391

Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role-based access control models. Computer, 29(2), 38–47. DOI: 10.1109/2.485845

Siau, K., & Wang, W. (2020). Artificial intelligence (AI) ethics: Ethics of AI and ethical AI. Journal of Database Management, 31(2), 74–87. DOI: 10.4018/JDM.2020040105

Sørheller, V. U., Høvik, E. J., Hustad, E., & Vassilakopoulou, P. (2018). Implementing cloud ERP solutions: A review of sociotechnical concerns. Procedia Computer Science, 138, 470–477. DOI: 10.1016/j.procs.2018.10.065

Thiebes, S., Lins, S., & Sunyaev, A. (2021). Trustworthy artificial intelligence. Electronic Markets, 31, 447–464. DOI: 10.1007/s12525-020-00441-4

Tikkinen-Piri, C., Rohunen, A., & Markkula, J. (2018). EU General Data Protection Regulation: Changes and implications for personal data collecting companies. Computer Law & Security Review, 34(1), 134–153. DOI: 10.1016/j.clsr.2017.05.015

Vilone, G., & Longo, L. (2021). Notions of explainability and evaluation approaches for explainable artificial intelligence. Information Fusion, 76, 89–106. DOI: 10.1016/j.inffus.2021.05.009