A JOINT WHITE PAPER

TEQ-IT × Xybern

AI Agents in SAP:

The Opportunity, the Risk, and the Governance Gap That Boards Are Not Being Told About

A practical guide for C-suite leaders navigating AI adoption on SAP S/4HANA, Ariba, and the broader SAP intelligent enterprise stack

May 2026

Introduction

SAP is no longer just an ERP system. For most of the organisations running it, SAP is the operational backbone of the business: the system that processes orders, manages suppliers, runs finance, controls inventory, and increasingly, drives decisions. And SAP is now deeply AI-enabled. SAP Joule, the AI copilot embedded across the SAP stack, can already execute tasks autonomously across S/4HANA, Ariba, SuccessFactors, and SAP Analytics Cloud. Third-party AI agents are being layered on top of SAP environments at pace. Custom-built automation is connecting AI models to SAP workflows in ways that procurement teams, finance functions, and operations leaders are embracing eagerly.

This is a genuinely exciting development. AI-enabled SAP systems can compress decision cycles, eliminate manual bottlenecks, reduce error rates, and surface insights that human

operators would miss. The business case for AI in SAP is not theoretical. It is being realised right now, in organisations that are moving quickly.

But there is a problem that almost nobody is talking about directly, and it is the subject of this white paper. When AI agents operate inside a SAP environment, they do not just surface recommendations. They take actions: creating purchase orders, approving invoices, updating master data, triggering payment runs, modifying supplier records, and delegating tasks to other agents. These are consequential, often irreversible actions. And in the vast majority of SAP deployments today, there is no mechanism that governs which actions an AI agent is permitted to take, records who authorised each action, and produces tamper-evident proof that the authorisation occurred.

That gap is not a theoretical compliance risk. It is a live operational exposure that auditors are beginning to identify, regulators are beginning to ask about, and boards are beginning to be held accountable for.

This paper is written for C-suite leaders at organisations running SAP or planning to adopt SAP S/4HANA. It is written jointly by TEQ-IT, a global SAP consulting partner with over a decade of implementation experience, and Xybern, which builds the authorisation infrastructure that AI-enabled enterprise systems need to operate safely and with full auditability. Together, we explain what the opportunity looks like, where the risk sits, and what organisations need to do about it before it becomes a board-level problem rather than a technology problem.

Why now?

SAP Joule and third-party AI agents are already active in production SAP environments. The 2027 SAP ECC end-of-maintenance deadline is accelerating S/4HANA migrations, and most organisations are embedding AI capabilities as they migrate. The decisions being made in migration projects today will determine whether AI governance is built in or bolted on later. Built in is significantly cheaper and safer.

SECTION 1 · AUTHORED BY TEQ-IT

The SAP AI Revolution: What Is Actually Happening

From ERP to Intelligent Enterprise

SAP's transformation over the past decade has been substantial. The move from SAP ECC to S/4HANA was more than a technical upgrade. S/4HANA is built on the in-memory HANA database, which processes data at speeds that change what is possible in real time. Reporting that once took hours runs in seconds. Closing processes that took days can be compressed to hours. And the architecture that underpins S/4HANA is the foundation on which SAP's AI capabilities are being built.

SAP describes its vision as the Intelligent Sustainable Enterprise: an organisation in which AI continuously analyses data, surfaces recommendations, automates routine decisions, and frees human workers for higher-value activity. This is not marketing language. It reflects the

genuine direction of SAP's product roadmap, and organisations that understand it are already building competitive advantage by adopting it.

SAP Joule: The AI Copilot That Is Already Executing Actions

SAP Joule is SAP's generative AI copilot, embedded across the SAP product portfolio. It is not a chatbot that answers questions. It is an AI system that can take actions inside SAP on behalf of users: creating and modifying records, running workflows, surfacing and acting on insights from across the data estate, and increasingly, operating autonomously between user interactions.

Joule is available across SAP S/4HANA, SAP Ariba, SAP SuccessFactors, SAP Analytics Cloud, and SAP Customer Experience. A user interacting with Joule in SAP Ariba can ask it to identify procurement anomalies and create a corrective action. A finance manager using Joule in S/4HANA can request that it close certain items and update forecasts. In SAP SuccessFactors, Joule can update employee records and trigger HR workflows. These are real capabilities available in production today.

The significance of this for organisations is that Joule is not a future capability being piloted. It is present in the SAP environments that organisations are deploying now, as part of S/4HANA migrations and new SAP cloud adoptions. The question is not whether to engage with it, but how to do so safely.

Third-Party AI Agents on SAP

Beyond Joule, a growing ecosystem of third-party AI agents is being integrated with SAP environments. Organisations are building custom agents using frameworks such as CrewAI, AutoGen, and LangGraph, connecting them to SAP via APIs and integration platforms. These agents are being used for demand forecasting, supplier risk monitoring, cash flow optimisation, anomaly detection in procurement, and dozens of other use cases where continuous AI analysis of SAP data creates business value.

These agents operate with varying degrees of autonomy. Some surface recommendations for human approval. Others are configured to act directly, creating records, triggering notifications, updating data, and in some cases delegating tasks to other agents in a chain. The more autonomous these agents become, the more consequential the governance question becomes.

The 2027 Migration Wave and the AI Adoption Moment

SAP will end mainstream maintenance for SAP ECC in 2027. This is not a soft deadline. After 2027, organisations remaining on ECC will not receive new functionality, regulatory updates will become uncertain, and security patches will no longer be available as standard. The practical implication is that every organisation running ECC must migrate to S/4HANA, move to a different ERP platform, or pay significantly for extended support.

Most organisations are choosing to migrate to S/4HANA. This is creating a significant wave of migration projects, and those projects are taking place at exactly the moment when SAP's AI capabilities are most feature-rich and most actively marketed. The result is that many organisations are adopting AI capabilities as part of their migration, embedding Joule and AI-enhanced workflows into their new S/4HANA environment from day one.

This is the right time to get AI governance right, because it is cheaper and simpler to build governance infrastructure into a migration project than to retrofit it into an established environment. But most migration projects are not doing this, because the governance requirement has not yet been made explicit to the organisations commissioning them.

What AI is already doing inside SAP environments today

- Creating and approving purchase orders in SAP Ariba without human review of individual transactions.
- Updating supplier master data based on AI analysis of risk signals.
- Triggering payment runs and modifying payment terms in S/4HANA.
- Modifying demand plans and inventory targets in SAP Integrated Business Planning.
- Delegating sub-tasks to other AI agents in multi-agent automation chains.
- Generating and distributing financial reports and management information.

SECTION 2 · AUTHORED BY TEQ-IT

The Business Case for AI in SAP: Real Benefits, Real Adoption

Why Organisations Are Moving Fast

The business case for AI in SAP is compelling and well-documented. Organisations that have deployed AI-enhanced SAP workflows are reporting material improvements across procurement efficiency, finance close speed, inventory accuracy, and operational decision quality. The benefits are not marginal. In procurement, AI-driven spend analysis and anomaly detection can identify savings that manual review would miss. In finance, AI-assisted period close can compress timelines from weeks to days. In supply chain, AI-enhanced demand sensing can reduce forecast error and buffer stock requirements significantly.

For C-suite leaders under pressure to demonstrate productivity gains, AI in SAP represents an unusually direct route from investment to measurable outcome. The capabilities are mature, the use cases are proven, and the integration with SAP's data estate means that AI has access to the operational data it needs to be genuinely useful, rather than working from incomplete or poorly structured inputs.

Procurement: Where the Gains Are Largest and the Risks Are Most Acute

SAP Ariba is one of the most AI-enhanced parts of the SAP portfolio, and procurement is the function where AI adoption is moving fastest. AI agents in Ariba can analyse millions of transactions to identify maverick spending, flag supplier risk, recommend contract consolidation opportunities, automate invoice matching, and generate guided buying recommendations that steer procurement behaviour without requiring manual category management.

The financial impact is significant. Organisations with AI-enhanced procurement typically see a measurable reduction in maverick spend, faster invoice processing cycles, and improved contract compliance. These are outcomes that CFOs can point to directly.

But procurement is also the function where the consequences of an unauthorised or incorrectly governed AI action are most immediately financial. An AI agent that creates a purchase order it was not authorised to create, approves an invoice outside its permitted scope, or modifies a supplier record without appropriate authority, creates a financial exposure that may not be visible until audit. And in a high-volume, AI-assisted procurement environment, by the time an audit identifies the exposure, hundreds or thousands of similar actions may have already occurred.

Finance: Speed and Accuracy at Scale

In finance functions, AI in S/4HANA is being used to automate journal entry creation, accelerate period close processes, improve cash flow forecasting, and surface anomalies in the general ledger that manual review would be unlikely to catch. These capabilities are genuinely transformative for finance teams that are under pressure to close faster, report more accurately, and provide more forward-looking insight to the business.

The governance dimension in finance is particularly sensitive. Finance processes are subject to internal controls frameworks, audit requirements, and in many jurisdictions regulatory reporting obligations. A finance AI that operates without an auditable record of its actions creates a gap in the internal controls framework that auditors will identify and that regulators may act on. The question of who authorised an AI to make a journal entry, under what control, with what approval, is not a hypothetical audit question. It is the kind of question that is already appearing in management letter points from external auditors who have reviewed AI-enabled finance environments.

Supply Chain: Intelligence at the Pace of Change

SAP Integrated Business Planning and the supply chain capabilities within S/4HANA are being enhanced with AI that can sense demand signals, optimise inventory positioning, and identify supply chain disruptions before they propagate. For organisations with complex supply chains, this represents a genuine step change in planning capability, moving from periodic planning cycles to continuous, AI-assisted sensing and response.

The governance question in supply chain AI is about the consequences of an incorrect or unauthorised AI decision at scale. A demand plan that has been modified by an AI agent acting outside its permitted scope can create inventory imbalances that take months to unwind. A procurement trigger that fires incorrectly because an AI agent interpreted a demand signal incorrectly can create committed spend that the organisation did not intend. Governance infrastructure does not prevent AI from being useful in supply chain. It ensures that when AI acts, it acts within defined boundaries, and that the evidence of those boundaries having been respected is available for review.

The core tension

The faster organisations move to adopt AI in SAP, the more consequential the governance gap becomes. Speed of adoption and depth of governance need to move together. The organisations that achieve both will capture the AI advantage without the audit exposure. The organisations that prioritise adoption without governance are accumulating a liability that will eventually become visible.

The Governance Gap: What Boards Are Not Being Told

The Question That Stops Most CIOs Cold

There is a question that, when asked to a CIO or CFO at an organisation running AI-enabled SAP, tends to produce a pause. The question is this:

> **The question**
> For any AI action taken inside your SAP environment in the last 30 days, can you tell me who authorised it, under what policy, and show me the proof?

The pause is not because the answer is complex. It is because, in almost every organisation running AI in SAP today, the honest answer is no. There is no mechanism that governs which actions an AI agent is permitted to take before it takes them. There is no versioned, auditable policy that defines what each AI agent can and cannot do. There is no tamper-evident record of the authorisation decision that preceded each AI action. And there is no cryptographic proof that any of this occurred.

This is not a failure of the organisations involved. It reflects the state of the market. SAP's AI capabilities have been developed and deployed faster than the governance frameworks required to manage them safely. Most SAP implementation projects, including most S/4HANA migration projects, do not include an AI governance workstream. The organisations commissioning those projects have not been told that one is needed.

Why Monitoring Is Not the Same as Governance

Many organisations that have thought about AI risk in their SAP environments have invested in monitoring: dashboards that show what AI agents are doing, alerts that fire when something unexpected happens, audit logs that record what occurred. Monitoring is valuable. But monitoring is not governance, and the distinction matters enormously.

Monitoring tells you what happened after it happened. Governance determines what is permitted to happen before it occurs. In a regulated environment, or in any environment where AI actions have financial or operational consequences, the difference between these two things is the difference between detecting a problem and preventing one.

An AI agent that creates an unauthorised purchase order in SAP Ariba will be visible in a monitoring dashboard after the purchase order has been created. The purchase order may already have triggered a commitment, notified a supplier, or initiated a goods receipt process. The damage is done before the monitoring alert fires. Governance would have intercepted the action before it executed, checked it against the authorised policy, and denied it if it was outside scope. No purchase order. No commitment. No audit finding.

The Audit Trail Problem

Audit trails in SAP record what happened inside the SAP system. Change logs show which records were modified and when. Workflow logs show which steps in a process were

executed. These are important and should be maintained. But SAP audit trails do not answer the governance question.

The governance question is not what the AI did. The governance question is whether the AI was authorised to do it, under what policy, with what scope, and where the cryptographic proof of that authorisation is. A SAP change log that shows an AI agent modified a supplier record does not tell an auditor whether that modification was within the agent's authorised scope. It does not show which version of the authorisation policy was in force at the time. It does not provide cryptographic proof that the authorisation decision was genuine and has not been altered. And it does not show the delegation chain if the action was taken by a sub-agent acting on behalf of a parent agent.

These are the questions that external auditors, internal audit functions, and regulators are beginning to ask about AI in enterprise systems. They are questions that SAP's own audit trail infrastructure is not designed to answer, because they are not questions about what SAP did. They are questions about the governance layer that should sit above SAP and govern what AI agents are permitted to ask SAP to do.

Real Scenarios Where the Gap Creates Exposure

Procurement authorisation drift

An organisation deploys an AI agent in SAP Ariba to automate purchase order creation for a defined category of indirect spend below a certain threshold. Over time, as the agent learns from data and the configuration drifts, the agent begins creating purchase orders for items outside its original category scope, at values above the intended threshold. Each individual transaction looks plausible. The drift is gradual. The monitoring dashboard does not flag it as an anomaly because each transaction is individually within tolerance. By the time the annual audit identifies the pattern, the organisation has months of committed spend that was not properly authorised.

A governance layer would have checked each proposed purchase order against the versioned policy before execution. Any action outside the defined category or above the defined threshold would have been denied before the purchase order was created. The audit finding would have been zero, because the actions that would have caused it never occurred.

Finance close AI and the internal controls gap

An organisation uses an AI agent to assist with the period close process in S/4HANA, automating journal entries for accruals and allocations based on rules defined at implementation. As the business changes, the rules are informally updated by a finance analyst with access to the agent configuration. The updates are not formally approved, the previous version of the rules is not retained, and the rationale for the changes is not documented. The external auditor asks for evidence of the internal control over the AI's journal entry creation. The organisation cannot provide it, because there is no versioned record of the rules under which the AI operated, no approval record for the rule changes, and no cryptographic proof that the rules in force at close were the rules that had been authorised.

This is an internal controls deficiency that an auditor is obliged to report. In a regulated industry, it may trigger a regulatory notification. In all industries, it creates reputational and governance exposure for the CFO and board.

Multi-agent delegation and the accountability gap

An organisation builds a supply chain automation using multiple AI agents: one to monitor demand signals, one to generate replenishment recommendations, one to create purchase

requisitions in SAP, and one to approve those requisitions within a defined value threshold. The agents communicate with each other, passing instructions and delegating actions. When an auditor asks who authorised a specific high-value purchase requisition, the organisation cannot trace the delegation chain. They can see that the requisition was created by the third agent, but they cannot show under what instruction from the second agent, within what scope delegated by the first, and with what cryptographic proof that the delegation was genuine and has not been reconstructed after the fact.

The governance gap in summary

- No pre-execution control that determines whether an AI action is permitted before it runs.
- No versioned, auditable authorisation policy linked to each AI agent.
- No cryptographic proof of the authorisation decision that preceded each action.
- No delegation chain record for multi-agent workflows.
- No tamper-evident record that cannot be altered or reconstructed after the fact.

SECTION 4 · AUTHORED BY XYBERN

How to Close the Gap: The Authorisation Layer for SAP AI

The governance gap described in Section 3 is real, but it is solvable. The solution is an authorisation layer that sits between AI agents and the SAP environment they interact with. Every action that an AI agent proposes passes through this layer before it reaches SAP. The layer checks the action against a versioned, auditable policy, verifies the identity of the agent, and returns a binary decision: authorise or deny. If authorised, the action proceeds to SAP. If denied, it is blocked before it executes. Every decision is recorded with cryptographic proof in an immutable Provenance Vault.

This is what Xybern provides. Not monitoring. Not guardrails. Not post-hoc audit logging. A mandatory authorisation pipeline, with cryptographic provenance, that governs every AI agent action before it occurs.

The Five-Stage Authorisation Pipeline

Xybern's authorisation layer operates as a five-stage pipeline that processes every AI action, without exception and without bypass:

- Intercept: Xybern sits between AI agents and the SAP environment. Every action, whether from SAP Joule, a third-party AI agent, or a custom automation, must pass through the authorisation pipeline before it reaches SAP. There is no route around it.
- Identify: Every AI agent carries a cryptographic identity. Xybern verifies exactly which agent is acting, under which role, with what trust level, and in what context, before any action is evaluated. Spoofing is structurally prevented.
- Authorise: The proposed action is checked against the current version of the authorisation policy. The policy is defined in code, versioned in source control, and

tested in shadow mode before deployment. It specifies precisely what each agent can and cannot do, in what context, within what value thresholds, and under what conditions.

- Decide: The layer returns a binary verdict, either authorise or deny. The decision is deterministic and traceable to the exact policy clause that governed it. There is no ambiguity, no scoring system, and no probabilistic threshold.
- Record: Every decision, including the action proposed, the agent identity, the policy version, the verdict, and the timestamp, is written to the Provenance Vault with a cryptographic signature and a hash chain linking it to the previous record. The record is immutable from the moment it is written.

How This Works in a SAP Environment

In a SAP S/4HANA or Ariba environment, the Xybern authorisation layer integrates between the AI agent layer and the SAP APIs or integration platform. When an AI agent proposes to create a purchase order in Ariba, the proposed action, including the supplier, the value, the category, and the requesting agent identity, is passed to Xybern before the Ariba API call is made. Xybern checks the action against the policy, verifies that the agent is authorised to create purchase orders for that supplier and category at that value, and returns a verdict. If the verdict is authorise, the Ariba API call proceeds. If it is deny, the API call is blocked, the denial is recorded in the Provenance Vault, and the appropriate escalation is triggered.

The same pipeline governs every other AI action across the SAP estate: journal entry creation in S/4HANA, supplier record updates in Ariba, demand plan modifications in Integrated Business Planning, report generation and distribution in SAP Analytics Cloud. Any action that an AI agent proposes to take inside or through SAP is subject to the authorisation pipeline.

This is framework-agnostic and model-agnostic. It works with SAP Joule, with custom agents built on any AI framework, and with third-party AI systems integrated via SAP's Business Technology Platform. It does not require changes to the SAP configuration or to the AI models themselves. It operates as a layer above both.

The Authorisation Policy: Versioned, Auditable, Linked to Controls

The authorisation policy is the document that defines what each AI agent is permitted to do. In Xybern's implementation, the policy is expressed in code, which means it is precise, testable, and versionable. It can be stored in source control alongside other configuration assets, reviewed and approved through a formal change management process, tested in shadow mode against live traffic before it goes into production, and linked directly to the organisation's internal controls framework.

This is significant for audit purposes. When an auditor asks what controls govern an AI agent's actions in SAP, the organisation can point to the authorisation policy as a documented, versioned, formally approved control. When the auditor asks what version of the policy was in force on a specific date, the version history in source control provides the answer. When the auditor asks whether the control operated as intended, the Provenance Vault records provide the evidence.

The policy is also the mechanism through which AI governance is integrated with the organisation's existing SAP authorisation framework. SAP's own authorisation objects, roles, and profiles define what human users are permitted to do inside SAP. The Xybern authorisation policy defines what AI agents are permitted to do. Together, they create a complete authorisation framework that covers both human and AI activity in the SAP environment.

The Provenance Vault: Cryptographic Proof for Audit

The Provenance Vault is the immutable record store where every authorisation decision is written. Its design is purpose-built for regulatory and audit requirements rather than operational monitoring.

Each record in the Provenance Vault is signed with an HMAC-SHA256 cryptographic signature that allows any reviewer to verify the record's authenticity. The records are linked in a SHA-256 hash chain: each record includes a hash of the previous record, so any alteration to any record in the sequence causes the chain to break and the alteration to become immediately detectable. Individual records can be disclosed to auditors using Merkle proofs, which allow a specific decision to be proven as genuine without exposing the full audit trail, which is exactly what litigation hold and regulatory review requirements demand.

The practical implication is that an organisation using Xybern can respond to any audit inquiry about an AI action in their SAP environment with a Provenance Vault record that is cryptographically verifiable, tamper-evident, and complete. The record shows what the agent proposed to do, whether it was authorised or denied, which version of the policy governed the decision, and when it occurred. No retrospective reconstruction. No reliance on mutable log files. Cryptographic proof.

Delegation Chains in Multi-Agent SAP Workflows

SAP environments frequently involve multi-agent workflows where one AI agent delegates tasks to others. A demand planning agent might delegate a procurement trigger to a purchasing agent, which delegates an approval to an approval agent. Each step in this chain must be governed, not just the final action.

Xybern's delegated authority chains ensure that when Agent A delegates to Agent B, the delegation is itself an authorised action, bounded in scope and time. Agent B cannot act beyond what was delegated to it. The scope of the delegation is cryptographically recorded. The full chain of authorisations, from the initiating agent through every delegated step to the final action in SAP, is available as a complete, verifiable record in the Provenance Vault.

This means that for any action taken by a multi-agent SAP workflow, an auditor can trace every step: which agent initiated the chain, what each agent was delegated to do, whether each delegation was within the authorised scope, and what the final action in SAP was and whether it was authorised. This is the level of accountability that regulators and auditors are beginning to require of AI-enabled enterprise systems.

Shadow Mode: Testing Before Enforcing

One of the practical concerns organisations have about implementing authorisation governance is the risk of disrupting established workflows. Xybern's shadow mode addresses this directly. When a new authorisation policy is being developed or when an existing policy is being updated, shadow mode allows the policy to run against live traffic without enforcing its decisions. The organisation can see exactly what the policy would have authorised and what it would have denied, without any impact on production operations.

This means that the policy can be calibrated against the actual behaviour of the AI agents in the environment before it is enforced, reducing the risk of legitimate actions being incorrectly denied. It also provides a documented evidence base for the policy design, showing what actions were observed in the environment and how the policy was designed to respond to them.

What the Xybern authorisation layer provides for SAP environments

- Pre-execution authorisation of every AI agent action, before it reaches SAP.
- Cryptographically verifiable, immutable Provenance Vault records of every decision.
- Versioned, auditable authorisation policies expressed in code, linked to the internal controls framework.
- Delegation chain records for multi-agent SAP workflows.
- Shadow mode testing for policy development and calibration.
- Human escalation queue for actions requiring operator review before proceeding.
- Framework-agnostic deployment: works with SAP Joule, custom agents, and any third-party AI on SAP.

SECTION 5 · AUTHORED BY TEQ-IT

Getting the SAP Foundation Right: Where Implementation Meets Governance

Governance infrastructure does not exist in isolation. It needs to be built on a well-implemented SAP environment with clean data, properly structured authorisations, and coherent process design. This is where TEQ-IT's role in the partnership becomes critical. The quality of the underlying SAP implementation determines how effectively governance can be applied, and a poorly structured SAP environment creates governance challenges that no authorisation layer can fully resolve.

Why SAP Implementation Quality Matters for AI Governance

The effectiveness of an AI agent operating in a SAP environment depends on the quality of the data and process structures it is working with. An AI agent that has access to clean, well-structured master data will make better decisions and be easier to govern than an agent operating in an environment where supplier records are duplicated, cost centres are poorly maintained, and approval hierarchies are inconsistently applied.

This is one of the arguments for investing in a well-structured S/4HANA migration rather than a technical lift-and-shift. An organisation that migrates to S/4HANA with its legacy data quality problems and process inconsistencies intact will find that AI agents amplify those problems rather than solve them. An agent that creates purchase orders based on ambiguous or duplicated supplier records will create governance problems that are much harder to resolve than the original data quality issue.

TEQ-IT's implementation methodology addresses this directly. Data cleansing and master data governance are foundational to a TEQ-IT implementation engagement, not afterthoughts. Process design is aligned with SAP best practice, which creates the clean process structures that AI agents can operate within predictably. And the authorisation framework within SAP, the roles, profiles, and segregation of duties controls, is designed as part of the implementation, not retrofitted after go-live.

Integrating AI Governance into the SAP Implementation Workstream

The most effective way to implement AI governance in a SAP environment is to treat it as a workstream within the SAP implementation project, not as a separate initiative that runs in parallel or follows the go-live. This is how TEQ-IT and Xybern work together.

In the design phase of an S/4HANA implementation, the AI use cases being considered are identified and scoped. The process flows that AI agents will participate in are documented. The actions that AI agents will be permitted to take are defined, along with the conditions and constraints that will govern them. This scoping work, done during implementation, becomes the input to the authorisation policy design.

In the build phase, the authorisation policies are developed in parallel with the AI agent configurations. Shadow mode testing is run against the development and quality environments to calibrate the policies against the actual behaviour of the agents. By the time the implementation reaches user acceptance testing, the authorisation layer is operational and the Provenance Vault is capturing records of every AI action in the test environment.

At go-live, the authorisation layer goes live with the rest of the implementation. AI governance is not a post-go-live workstream. It is part of the delivered system.

The Post-Migration Landscape: Ongoing Governance as Operations Evolve

An S/4HANA implementation is not a static endpoint. The SAP environment will evolve after go-live: new processes will be added, existing processes will be modified, AI agent configurations will be updated, and new AI capabilities from SAP and third parties will be adopted. Each of these changes has governance implications.

TEQ-IT's support and managed services model includes AI governance as an ongoing operational concern. When process changes are made in SAP, the authorisation policies are reviewed and updated as part of the change management process. When new AI capabilities are adopted, the policy scope is extended to cover them before they are moved to production. When the SAP environment is upgraded or extended, the authorisation layer is tested as part of the upgrade validation.

This ongoing governance model ensures that the authorisation layer remains aligned with the evolving SAP environment and that the Provenance Vault continues to provide complete coverage of AI activity as the system matures.

Regulatory and Compliance Considerations

For organisations in regulated industries, the governance of AI in SAP is increasingly a regulatory matter rather than just a good practice recommendation. In financial services, regulators including the PRA and FCA are publishing expectations around the governance of AI systems used in regulated activities, including systems that interface with financial data and transaction processing. In healthcare, MHRA and EU MDR requirements extend to AI systems that interact with clinical or patient-related data. In any sector subject to SOX compliance, the governance of AI systems that touch financial reporting processes is directly relevant to the internal controls framework.

TEQ-IT's experience across multiple regulated sectors, combined with Xybern's governance infrastructure, means that organisations in these sectors have a partnership that understands both the regulatory expectations and the technical implementation required to meet them. The Provenance Vault records, the versioned authorisation policies, and the shadow mode testing evidence are all designed to be presented to regulators and auditors as evidence of effective AI governance.

What TEQ-IT brings to AI governance in SAP

- Clean, best-practice S/4HANA and Ariba implementations that create the foundation for effective AI governance.
- Master data governance and data quality programmes that ensure AI agents work with reliable inputs.
- Integration of AI governance into the implementation workstream, not as a post-go-live afterthought.
- Ongoing managed services and change management that keep governance aligned as the environment evolves.
- Sector-specific expertise in regulated industries where AI governance has regulatory as well as operational significance.

SECTION 6 · AUTHORED JOINTLY BY TEQ-IT AND XYBERN

A Practical Roadmap for C-Suite Leaders

The challenge described in this paper is significant, but it is navigable. Organisations that approach AI in SAP with both ambition and governance discipline will capture the operational benefits of AI without accumulating the audit and compliance exposure that ungoverned AI creates. This section sets out a practical roadmap: what to do at each stage, how TEQ-IT and Xybern work together in practice, and what questions to ask if you are evaluating your current position.

Stage 1: Assess Your Current AI Exposure

The starting point for most organisations is an honest assessment of the AI activity that is already occurring in their SAP environment. This is often more extensive than leadership teams realise. SAP Joule may be active in environments that were configured to include it as part of a cloud subscription. Third-party integrations may be executing automated actions that meet the definition of AI agent behaviour even if they were not described as AI when they were implemented. Custom automation built on older RPA or workflow platforms may have been extended with AI capabilities without a formal governance review.

The assessment should identify every process in the SAP environment where AI is involved in generating, approving, or executing actions. For each such process, the assessment should determine whether there is a documented, versioned, formally approved control governing what the AI is permitted to do, and whether there is an auditable record of the AI's actions that would satisfy an external auditor.

For most organisations, this assessment will reveal a governance gap that is larger than expected. The output of the assessment is the input to the governance programme.

Stage 2: Prioritise by Risk and Value

Not every AI action in a SAP environment carries the same risk or creates the same value. The governance programme should be prioritised based on a clear-eyed assessment of where the risk is highest and where the value of the AI is greatest.

Risk is typically highest where AI actions are financially consequential, directly irreversible, or subject to regulatory scrutiny. Purchase order creation, invoice approval, journal entry creation, payment processing, and supplier master data management are all high-priority areas for governance. Risk is lower where AI actions are informational, where a human reviews and approves the AI recommendation before it is acted on, or where the consequences of an incorrect action are minor and easily reversed.

Value is typically highest where AI is operating at a volume or speed that human oversight cannot match: high-volume procurement transactions, continuous monitoring of large data sets, real-time demand sensing, and automated period close activities. These are also the areas where governance infrastructure is most important, because the volume of AI actions makes manual review impractical and the need for automated, policy-based control most acute.

Stage 3: Build Governance Into What Is Being Built

For organisations currently in an S/4HANA migration or planning one, the most important action is to ensure that AI governance is included in the programme scope before the implementation begins. Adding governance infrastructure to an established implementation is significantly more expensive than including it from the start.

The scope addition required is not large in the context of an S/4HANA programme. It involves including an AI use case scoping exercise in the design phase, incorporating policy design and shadow mode testing into the build and test phases, and deploying the authorisation layer as part of the go-live package. TEQ-IT and Xybern can provide a joint scoping assessment that identifies the governance requirements for a specific implementation and defines the workstream structure needed to meet them.

Stage 4: Retrofit Governance to Existing AI-Enabled Environments

For organisations that already have AI active in their SAP environment and need to address the governance gap retroactively, the programme is different but equally achievable. Xybern deploys in under one week per workflow, without requiring changes to the SAP configuration or to the AI models. The deployment sequence begins with the highest-risk workflows, identified in the Stage 1 assessment, and extends coverage progressively across the AI estate.

Shadow mode is used for the initial calibration of each policy, allowing the organisation to understand the current behaviour of AI agents in each workflow before enforcement begins. This typically reveals useful information about scope drift, threshold violations, and delegation patterns that have developed since the AI was first deployed. The shadow mode findings inform the final policy design, and enforcement begins once the policy has been validated.

How TEQ-IT and Xybern Work Together

TEQ-IT leads on SAP strategy, implementation, process design, data governance, and the SAP-specific aspects of the governance programme. This includes the design of the SAP authorisation framework, the identification of AI use cases within the SAP environment, and the integration of governance requirements into the implementation methodology.

Xybern leads on the authorisation layer implementation, policy design, Provenance Vault configuration, and the translation of governance requirements into executable authorisation policies. Xybern also provides the ongoing management of the authorisation layer as the SAP environment evolves.

The integration point between the two is the risk and controls framework. TEQ-IT defines the process controls required within the SAP environment, including the controls that should govern AI actions. Xybern implements those controls as authorisation policies and generates the Provenance Vault evidence that the controls are operating as intended. The result is a governance programme that is fully integrated with the SAP implementation, rather than a separate layer that needs to be reconciled with it.

Questions to Ask to Assess Your Current Position

The following questions will help you assess the current state of AI governance in your SAP environment and identify the most important areas for action:

- Can you produce a complete list of every AI agent, SAP Joule workflow, or automated process that takes actions inside your SAP environment?
- For each AI action in your SAP environment, is there a documented, versioned, formally approved control that defines what the AI is permitted to do?
- If an external auditor asked to see the evidence that a specific AI action on a specific date was within the AI's authorised scope, could you produce it?
- Is there a cryptographically verifiable, tamper-evident record of every AI action in your SAP environment, including the authorisation decision that preceded it?
- For multi-agent SAP workflows, can you trace the complete delegation chain for any action, from the initiating agent to the final SAP action?
- Does your S/4HANA migration programme include an AI governance workstream, or is AI governance being treated as a post-go-live consideration?
- Has your external auditor or internal audit function reviewed the governance of AI in your SAP environment in the last twelve months?

If the answer to any of these questions is no, or if the answer is uncertain, the governance gap is present and the programme to address it should begin.

The three actions to take first

- Conduct an AI exposure assessment to identify every AI agent and automated process active in your SAP environment.
- Prioritise the highest-risk workflows, typically in procurement, finance, and supply chain, for immediate governance coverage.
- If you are in an S/4HANA migration, ensure that AI governance is included in the programme scope before the design phase closes.

Conclusion

The integration of AI into SAP environments is one of the most significant operational developments facing enterprise organisations in 2026. The benefits are real, proven, and accessible. SAP Joule, third-party AI agents, and custom automation are delivering measurable improvements in procurement efficiency, finance close speed, supply chain intelligence, and operational decision quality. The organisations embracing these capabilities are building competitive advantage.

But the governance gap that exists in most AI-enabled SAP environments is equally real, and its consequences are becoming more visible. Auditors are asking questions about AI governance that organisations cannot answer. Regulators are developing expectations that will become requirements. Boards are beginning to understand that ungoverned AI in an enterprise system is not a technology risk but a governance risk, and governance risk is their responsibility.

The window to address this proactively is now. Organisations in the middle of an S/4HANA migration have the clearest opportunity: governance infrastructure can be built into the implementation at a fraction of the cost of retrofitting it. Organisations with AI already active in their SAP environment need to assess their exposure and begin the governance programme without waiting for an audit finding to make it urgent.

TEQ-IT and Xybern provide an integrated programme that addresses both the SAP implementation quality and the AI governance infrastructure required to operate safely. The SAP expertise and the authorisation layer are designed to work together, because the organisations that will capture the full value of AI in SAP are those that build it right from the start.

Key takeaways

- AI agents in SAP are already taking consequential, often irreversible actions in production environments.
- The governance gap, the absence of pre-execution authorisation, versioned policy, and cryptographic proof, is present in the vast majority of AI-enabled SAP deployments.
- Monitoring and SAP audit trails do not close the governance gap. An authorisation layer that governs what AI can do before it does it is required.
- The 2027 ECC deadline is creating an S/4HANA migration wave. Organisations that include AI governance in their migration programme will deliver it at significantly lower cost and risk.
- TEQ-IT and Xybern provide an integrated programme: SAP implementation quality from TEQ-IT, authorisation infrastructure and cryptographic provenance from Xybern.
- The time to act is before the audit finding, not after.

About TEQ-IT and Xybern

TEQ-IT

TEQ-IT is a global SAP consulting partner specialising in the implementation, roll-out, and ongoing support of SAP solutions for organisations of all sizes, from growing mid-market businesses to large international enterprises. With over ten years of SAP experience and a team of expert consultants spanning SAP S/4HANA, SAP Business One, SAP Analytics Cloud, SAP Ariba, and the full SAP Customer Experience suite, TEQ-IT guides clients through the full lifecycle of SAP adoption.

TEQ-IT's services span strategy and solution design, full-cycle implementation, global roll-out programmes, managed services, and web and software development. The firm works with clients across multiple sectors and geographies, with a particular focus on helping organisations navigate complex SAP migrations and unlock the business value of SAP's intelligent enterprise capabilities.

TEQ-IT's approach combines deep technical SAP expertise with a business-first perspective: the goal of every engagement is not a successful go-live but a sustainable improvement in the client's operational performance and competitive position.

Contact TEQ-IT

Website: teq-it.com
Email: info@teq-it.com
Phone: +44 7466 836000
Address: 183, Tolworth Rise North, Surbiton, KT5 9ES, United Kingdom

Xybern

Xybern builds the authorisation layer for AI agents. In enterprise environments where AI agents trigger payments, query production databases, execute procurement actions, and delegate work to other agents, Xybern intercepts every action before it executes and returns a clear authorise-or-deny verdict. Every decision is recorded in the Provenance Vault with cryptographic signatures and hash chains, creating tamper-evident, auditable proof that is ready for regulatory review, external audit, and internal controls assessment.

Xybern is not a monitoring platform, an observability layer, or a guardrails tool. It is a mandatory authorisation pipeline: every AI agent action, governed before it runs. The Provenance Vault generates the cryptographically verifiable records that EU AI Act compliance, SOX internal controls, FCA and PRA AI governance expectations, and enterprise audit requirements demand.

Xybern is framework-agnostic and model-agnostic, integrating with SAP Joule, CrewAI, AutoGen, LangGraph, and any custom agent architecture, without requiring changes to the underlying AI systems. It deploys in under one week per workflow and is backed by NVIDIA Inception and AWS for Startups.

Contact Xybern

Website: xybern.com
Email: info@xybern.com
Pilot enquiries: xybern.com/enterprise

Ready to govern your AI in SAP?

Speak to TEQ-IT and Xybern to discuss your current SAP environment, your AI adoption plans, and how to close the governance gap before your next audit. The earlier you engage, the more straightforward the programme.