Governing AI Access in SAP Landscapes: S/4HANA, Business AI, and Hybrid Risk

SAP was already one of the hardest parts of your control environment before AI. You are managing S/4HANA timelines, RISE contracts, legacy ECC and BW, and a role and SoD model that stretches across finance, logistics, HR, and procurement. Now SAP Business AI, BTP-hosted models, and external agents are starting to read, recommend, and even act across that same landscape. The real risk is not "AI in SAP" as a concept; it is AI quietly inheriting SAP authorizations and cross-system access paths that were never designed for autonomous, always-on agents.

As AI gets embedded into S/4HANA, extended through SAP BTP, and connected to non-SAP systems, the control problem becomes architectural. Decisions are proposed in one place, executed in another, and audited—if at all—in a third system. Governing AI in SAP means treating agents and copilots as first-class identities in your SAP access model, not as invisible layers sitting on top of existing users and RFC/service accounts.

This article builds on AI Governance: When AI Becomes an Identity and Top 5 AI Access Risks for CISOs and How AI Governance Closes the Gaps, and shows what good AI governance looks like in SAP landscapes. For the SAP-specific access-governance view, see also Automate SAP Access Governance for Better Compliance and Segregation of Duties for SAP | SafePaaS Solutions.

How AI actually shows up in SAP

In a typical SAP estate (S/4HANA plus surrounding systems), AI now appears in three main patterns.

1. Embedded SAP Business AI inside S/4HANA and cloud apps

SAP delivers AI-driven recommendations and copilots directly inside finance, procurement, and supply chain transactions—for example, Cash Application in S/4HANA, AI-assisted forecasting, and Joule, SAP's enterprise copilot, which now spans S/4HANA Cloud, SuccessFactors, Ariba, and BTP with thousands of skills. These capabilities run under the same user and role concepts you already use—business roles, authorization objects, Fiori catalogs—so any over-privileged role instantly becomes an over-privileged AI surface. Joule, for instance, inherits the calling user's full authorization profile: if a business role grants broad change authority, Joule can act on it.

If a credit controller role already has broad change authority, and you let an embedded assistant inherit that role, you have effectively given an AI agent the same power to influence credit limits, dunning, and disputes without separately governing that identity.

2. Side-by-side AI on SAP BTP and integrated platforms

Custom models and agents built on SAP BTP, SAP AI Core, or external AI platforms consume SAP data via OData services, CDS views, and APIs, then push decisions back into S/4HANA, Ariba, SuccessFactors, or Fieldglass. In many cases, they run under shared technical users, communication users, or RFC/service accounts that cut across multiple systems and modules.

A single account might be able to read FI/CO documents in S/4HANA, update supplier data in Ariba, and touch HR data in SuccessFactors. AI agents built on top of those accounts inherit that combined authority, yet they rarely appear as distinct identities in SAP GRC or IAG.

3. Hybrid and non-SAP AI around the SAP core

Data replicated from S/4HANA into BW/4HANA, data lakes, or third-party analytics platforms is increasingly fed into AI tools or enterprise agents. These agents may influence pricing, approvals, and forecasts, while the execution authority still sits in SAP workflows—creating a split between where decisions are made and where they are enforced.

When AI operates on SAP-derived data outside SAP, its outputs still drive actions in SAP (for example, automated PO approvals, forecast-driven MRP changes, or cash-management decisions), but the AI identities and access paths are often invisible to SAP security teams.

Across all three patterns, the underlying problem is the same: AI agents inherit SAP roles and cross-system access paths that were built for people and traditional integrations, not for autonomous, scaled-out decision-makers. For the cross-platform picture, see Federated Governance for AI Identities: Closing the 92% Visibility Gap and the SAP-specific SoD guidance in SAP Authorisation Best Practices: Avoiding SoD Conflicts.

What “good” looks like in SAP

For SAP estates, best-in-class AI governance looks like modern SAP access governance with four extensions.

1. AI agents are first-class SAP identities

Each SAP-adjacent copilot, BTP agent, or external AI that can influence SAP data or transactions is represented as an identity—technical user, communication user, or managed account—that appears in SAP Access Control / SAP Cloud IAG and in your central IGA platform.

Each identity has a named owner, business purpose, and risk classification, just like a high-risk SAP user.

2. AI business roles are separate from human roles

AI use cases consume AI-specific business roles composed of fine-grained authorization objects and Fiori catalogs, rather than inheriting existing “power user” roles.

These roles are designed with SoD and least-privilege in mind from the start and are blocked from being assigned to human users.

3. SAP SoD and critical access rules apply to AI too

Your SAP SoD ruleset (for example, “Create Vendor” vs “Release Payment,” “Create PO” vs “Post Goods Receipt”) explicitly includes AI roles and technical users.

AI identities are subject to the same access risk analysis, mitigations, and periodic certifications as high-risk human accounts. SafePaaS’s cross-system rule sets, described in Segregation of Duties for SAP | SafePaaS Solutions, already incorporate S/4HANA processes, Ariba, Concur, and SuccessFactors.

4. SAP GRC / Cloud IAG and enterprise IGA form one control plane

SAP Access Control or SAP Cloud Identity Access Governance handle SAP-specific risk analysis and provisioning, while an enterprise IGA platform orchestrates policy-based access, lifecycle, and certification across SAP and non-SAP systems.

AI identities are normal citizens in that control plane, not one-off exceptions managed in code. SafePaaS provides this federated control layer across SAP, Oracle, and SaaS—see How to Govern AI Access to ERP and Financial Systems and Access Governance and Risk Management.

JML for SAP AI: an “AI Goods Receipt Assistant” example

To make this concrete, imagine a simple but high-impact use case: an “AI Goods Receipt Assistant” for S/4HANA that helps warehouse teams match inbound deliveries against POs, flag discrepancies, and propose goods receipts—but cannot post receipts or change material master data.

Joiner: onboarding the AI Goods Receipt Assistant

Intake the use case in SAP terms

You document the process (“GR/IR matching and goods receipt support”), in-scope systems (S/4HANA, possibly EWM), company codes, plants, and document types. You define data needs (PO line items, inbound deliveries, historical discrepancies) and permitted actions (read PO and delivery data, propose goods receipts, no posting). You tag regulatory scope if applicable (for example, SOX for inventory and COGS). The CISO Toolkit for AI Identity & Access Governance helps standardize this intake and risk assessment.

Design AI-specific SAP roles and authorizations

Instead of reusing warehouse or MM roles, you design a dedicated AI business role, such as Z_AI_GR_ASSISTANT, that includes only:

- Display authorizations for relevant MM, MM-IM, and LE objects (POs, deliveries, material documents) in defined plants and company codes.
- Authorizations to create proposed goods receipt documents in a staging area or workflow, not to post them.
- No access to vendor master maintenance, price conditions, or configuration, and no direct posting or reversal rights.

This role is built with SAP Access Control / Cloud IAG guidance and tested against your SoD ruleset to ensure it is “clean by design,” following the same patterns outlined in How Experienced SAP Teams Check Segregation of Duties (SoD).

Create a governed SAP technical identity for the agent

You provision a dedicated technical user or communication user for the AI agent in S/4HANA (and, if relevant, in BTP). That identity:

- Is tagged as “Non-human / AI” in your identity store.
- Holds only Z_AI_GR_ASSISTANT and any required communication roles.
- Is linked to a business owner (for example, Head of Logistics) and a technical owner (for example, SAP platform lead).

Run SoD and risk analysis before go-live

Before the AI agent is connected, you run SAP access risk analysis against the proposed roles. You confirm there are no SoD conflicts (for example, the agent cannot both propose receipts and release payments, or both create POs and post receipts) and that critical access does not exceed what is justified in the intake.

Provision through a policy-driven workflow

The request to create and enable the AI GR Assistant identity flows through your IGA / SAP GRC or Cloud IAG workflows:

- AI use-case owner submits the request.
- Risk and SoD checks are performed automatically.
- Approvals are collected from the logistics process owner and risk or internal audit.
- S/4HANA and BTP provisioning occurs automatically on approval.

All of this generates an audit trail—who requested, who approved, what was granted, and what risks were analyzed.

Mover: changing AI scope in SAP

Over time, your AI Goods Receipt Assistant might expand from a single plant to a region, or from one company code to several.

Scope change triggers a new request
Any proposed expansion in plant, company code, document type, or action (for example, from “propose only” to “post low-value receipts”) must go through a mover workflow, not a quiet change to the underlying role.

Re-run SoD and update the authorization model
If you are only expanding data scope, you might add new organizational values to existing authorization objects and re-run risk analysis. If you are adding new capabilities—such as limited posting rights—you design a separate Z_AI_GR_POST_LOW_VALUE role with strict constraints (for example, value thresholds, document types, specific plants) and test it against SoD and critical access rules.

Escalate approvals with risk
Changes that increase AI posting capability or expand to new critical plants and company codes should automatically route to higher-level approvers—CFO, regional controller, or an AI governance committee—because they change financial reporting and operational risk.

Continuously optimize authorizations
Using access analytics and usage data, you track which authorizations the AI actually uses.

Over time, you remove unused authorizations and refine Z_AI roles, converging toward true least-privilege for the agent.

Leaver: retiring AI identities and integrations in SAP

AI pilots end, agents are replaced, and integrations move to new platforms. Without lifecycle controls, SAP estates accumulate “zombie” technical users with lingering rights.

A clean leaver process for AI in SAP includes:

- **Event-based offboarding** – When a project ends, a contract expires, or an AI identity has been inactive for a defined period, it is flagged as a leaver in your IGA / SAP IAG system.
- **Revoking SAP and BTP access** – The AI technical user is locked and then deleted or end-dated; associated roles and authorizations are removed. Any BTP destinations, OAuth clients, or API keys linked to that identity are revoked.
- **Preserving logs and evidence** – You retain identity records, approval histories, and key activity logs for audit and forensic purposes, but eliminate the ability to log in, call services, or trigger jobs.

The outcome: AI identities follow the same Joiner–Mover–Leaver discipline as human users, and you can show a closed loop from creation to deprovisioning.

A practical SAP AI governance checklist

To make this actionable for SAP program owners, CISOs, and CFOs, you can use a concise checklist over the next 90 days:

- **Inventory AI-adjacent and non-human identities across SAP**
  Identify technical and communication users, RFC/service accounts, bots, and early AI agents in S/4HANA, ECC, BTP, SuccessFactors, Ariba, and Concur. SafePaaS’s SAP solutions overview at Automate SAP Access Governance for Better Compliance is a good reference for what to include.
- **Map their roles, authorizations, and SoD risks**
  Use SAP Access Control, SAP Cloud IAG, or a partner solution to run access risk analysis and identify SoD and critical-access issues for those identities. For patterns and pitfalls, see How Experienced SAP Teams Check Segregation of Duties (SoD).
- **Define AI-specific business roles in SAP**
  Create a small number of AI-only roles (for example, Z_AI_AP_ANALYST, Z_AI_GR_ASSISTANT, Z_AI_CASH_APP) with clearly scoped authorization objects and SoD-clean design, and block assignment to human users.

- **Bring AI into your SAP and enterprise JML workflows**
  Ensure any new AI agents that can read or act on SAP data are created via standard access request workflows with intake, risk assessment, and approvals—never as ad-hoc technical users.
- **Automate provisioning and deprovisioning for AI identities**
  Integrate SAP IAG / Access Control with your IGA platform so AI identities are provisioned and revoked automatically across SAP and non-SAP systems based on lifecycle events. SafePaaS's federated governance model in Access Governance and Risk Management shows what this looks like across platforms.
- **Include AI identities in SAP access reviews**
  Make AI and non-human identities part of periodic SAP user access reviews. Ask business owners to attest to necessity, scope, and risk for each AI identity.
- **Monitor AI behaviour across SAP systems**
  Feed SAP logs and AI agent activity into your SIEM and analytics platforms, with clear markers for AI vs human actions, and define policies for when to alert, quarantine, or revoke access.

Handled this way, AI inside SAP stops being an opaque add-on and becomes another class of identity you manage with the same discipline as your most critical SAP users. You can move faster on SAP Business AI and BTP-based innovation while giving your board and regulators something they rarely get in this space: a clear, evidence-backed story about which agents touch your SAP data and processes, what they can do, and how that access is governed over time.

Next step: book a SAP AI access governance session with SafePaaS to see where AI and non-human identities already touch S/4HANA, Ariba, and SuccessFactors, and how a federated control plane can bring them under policy and evidence.

Talk to SafePaaS about SAP AI access governance

.