



ALPHAOAK

White paper

GOVERNING AI AGENTS IN SAP BUSINESS OPERATIONS

Audit traceability in the agentic SAP landscape

A business-control framework for accountable human
and autonomous agent transactions

Tony Buffington

Published by AlphaOak · October 2026

**BUSINESS
PROCESS**

Define tasks and
decisions

AUTHORITY

Assign human and
agent rights

EVIDENCE

Preserve the audit
chain

CONTENTS

What this paper argues, in ten parts

HOW TO READ THIS PAPER

A white paper for the people who will have to explain an AI agent's transaction after the fact. Parts 01 to 05 set out the evidence problem; Parts 06 to 09 turn it into a control design; Part 10 describes how AlphaOak helps. Sources are listed at the back.

01 The audit problem changes when an agent becomes an actor

A technically valid SAP posting can hide an invisible decision path. The audit has to answer more than who posted it.

02 Two operating models require different evidence

Supervised agents need proof of what the human saw and decided. Autonomous agents move accountability upstream into policy.

03 Identity and transaction authority form the control boundary

An agent's effective authority is the overlap of four boundaries, never a copy of the requester's broad role.

04 The audit trail is a chain of correlated evidence

No single log tells the whole story. A shared correlation ID has to connect eight stages from initiation to outcome.

05 What a defensible audit record should contain

Eleven evidence categories, presented as one business event an auditor can read without rebuilding thousands of system events.

06 A business-process control model for agent transactions

Auditability begins before the agent is built: every task gets an actor, an authority boundary and an evidence obligation.

07 Audit walkthrough: maintenance-risk intervention

Seven steps in an SAP EAM scenario show how process design and technical records work together.

08 Controls across the agent lifecycle

Eight lifecycle stages, each with a control objective and evidence, and testing that deliberately exercises the failure paths.

09 Implementation principles

Eight principles, starting with the business process rather than the model, and ending with outcomes rather than uptime.

10 AlphaOak's role in creating audit-ready agentic processes

Six areas of work that make the process, authority model, human-agent division of work and audit evidence explicit.

— Conclusion, practical starting point and references

EXECUTIVE SUMMARY

The audit question is no longer only who changed the record

AI agents are becoming actors inside enterprise processes. Some recommend actions for a human to approve; others execute predefined business tasks autonomously. In both cases, the central audit question is no longer only who changed an SAP record. The organization must be able to reconstruct the chain from business intent, through human and agent decisions, to the resulting transaction and outcome.

SAP provides important technical mechanisms for identity, policy enforcement, agent inventory, observability, security logging and business-object change history. These mechanisms do not automatically form a complete business audit. Clients must define the business process, authority model, control points and evidence requirements that connect the records into a defensible account of what occurred.

THE PROPOSITION, STATED PLAINLY

An auditable agent transaction must preserve business purpose, human accountability, agent identity, authorization, decision lineage, tool execution and the final SAP business-object change.

What this paper establishes

- A practical distinction between autonomous and human-supervised agent activity.
- The minimum evidence required to trace an AI-assisted SAP transaction.
- A control framework spanning process, identity, authorization, execution and outcome.
- A target audit record that correlates evidence across SAP and supporting platforms.
- AlphaOak's role in designing, documenting, implementing and validating this business capability.

Who it is for

Business-process owners, internal audit, enterprise architects, SAP security teams, AI governance leaders, risk and compliance functions, and program sponsors responsible for deploying AI into business operations.

PART 01

The audit problem changes when an agent becomes an actor

Traditional SAP audits are organized around recognizable actors and records: a user ID, a transaction or application, an authorization role, a document number and a change log. Agentic systems introduce another decision-making and execution layer between business intent and the system of record.

An agent can interpret a goal, select tools, retrieve current data, evaluate alternatives, invoke another agent and initiate an API call. The final SAP transaction may be complete and technically valid while the broader decision path remains invisible unless the design deliberately preserves it.

The audit must answer more than “who posted it?”

AUDIT QUESTION	EVIDENCE REQUIRED
What initiated the action?	Human request, schedule, business event or another agent.
Who or what acted?	Human identity, agent identity, runtime or integration identity.
What authority applied?	User permissions, agent permissions, tool scope and policy decision.
What information was used?	Business objects, data sources, versions and sensitive-data access.
How was the decision reached?	Rule, prompt, model, agent plan, recommendation and exception path.
Where was human judgment applied?	Approval, rejection, modification, escalation or override.
What changed?	SAP document, object, field-level before/after values and result.
Was value delivered safely?	Outcome, exceptions, compliance, performance and follow-up.

KEY IMPLICATION

A technically logged transaction is not necessarily an auditable agent decision. Auditability requires correlation across the entire business event.

PART 02

Two operating models require different evidence

2.1 Human-supervised agents

A non-autonomous or human-supervised agent gathers information, reasons or recommends, but a person remains in the decision or execution path. The strength of the audit depends on preserving the separation between the agent's proposal, the human's decision and the transaction that was actually posted.

- Record the initiating human and the context of the request.
- Preserve the recommendation presented for review.
- Record whether the person approved, rejected or modified the recommendation.
- Capture the identity and authority used for the final action.
- Prove that the executed SAP transaction matched the approved instruction.

2.2 Autonomous agents

An autonomous agent performs approved tasks without transaction-by-transaction human approval. Accountability does not disappear; it moves upstream into the approved policy, operating boundary and ownership model.

- Identify the business owner who authorized the autonomous purpose.
- Define exactly which triggers, objects, values and actions are permitted.
- Assign a dedicated agent identity and minimum required authority.
- Record every execution, decision, tool call, transaction result and exception.
- Monitor behavior against approved process and policy boundaries.
- Require escalation when confidence, value, risk or exception thresholds are exceeded.

CONTROL CHARACTERISTIC	HUMAN-SUPERVISED	AUTONOMOUS
Decision authority	Named human approves or modifies	Approved policy authorizes bounded execution
Primary accountability evidence	Human approval and final transaction	Policy, ownership, trigger and execution trace
Exception handling	Human resolves before execution	Agent stops, retries or escalates under defined rules
Audit emphasis	Proposal → approval → posting	Policy → trigger → decision → posting

PART 03

Identity and transaction authority form the control boundary

Agents should access SAP through approved tools and APIs under authenticated, constrained identities. They should not receive unrestricted access simply because the initiating person has broad permissions.

EFFECTIVE AUTHORITY



An agent acting for a broadly authorized user still gets only the overlap of these four boundaries.

Identities that should survive the transaction

IDENTITY	AUDIT PURPOSE
Initiating human	Who requested the activity or supplied the business intent.
Approving human	Who accepted, rejected or modified a consequential recommendation.
Agent identity	Which registered agent planned, recommended or executed the work.
Tool or service identity	Which API, MCP server or integration component called SAP.
Target-system identity	Which identity the SAP application recorded against the transaction.
Business owner	Who owns the approved purpose, policy and risk decision.
Technical owner	Who owns deployment, configuration, monitoring and remediation.

IMPORTANT DISTINCTION

Audit records should retain identity attributes, role and scope references, token identifiers or policy decisions where appropriate. They should not retain passwords, client secrets or complete bearer tokens. Secrets belong in managed credential stores and must remain outside business audit evidence.


SAP's current agent architecture describes both agents operating in a user's context and autonomous agents operating through dedicated identities. For user-context activity, SAP states that effective permission at runtime is the intersection of user and agent permissions, enforced by the Agent Gateway at every communication hop. [\[source\]](#)

PART 04

The audit trail is a chain of correlated evidence

No single log is likely to contain the complete business story. Evidence is distributed across the user experience, agent runtime, identity and policy services, tools and integrations, and the SAP system of record. A shared correlation identifier must connect them.

END-TO-END EVIDENCE CHAIN



■ Human control point
 ■ SAP system of record

TRACEABILITY RULE

Every material agent action should carry a correlation ID from initiation through the resulting SAP document and operational evidence.

Potential SAP evidence sources

Depending on the selected products and deployment pattern, relevant records may include Joule interaction and agent traces, Agent Gateway policy and tool events, SAP Cloud Identity Services events, MCP or Integration Suite logs, SAP BTP audit log records, S/4HANA security audit logs, change documents, workflow approvals, Read Access Logging, and SAP Signavio process or agent-behavior mining evidence.

PART 05

What a defensible audit record should contain

The audit record should present the material business evidence while preserving links to the underlying technical logs. It should be understandable to an auditor without requiring the auditor to reconstruct thousands of unrelated system events.

EVIDENCE CATEGORY	MINIMUM CONTENT
Business context	Process, use case, purpose, risk classification and accountable owner.
Trigger	Request, event or schedule; initiating person or system; timestamp.
Agent	Stable agent ID, version, autonomy level, owner and approved capability.
Authorization	Human context, agent identity, policy result, roles/scopes and tool boundary.
Data accessed	Source, business objects, relevant record keys and sensitive-data indication.
Decision lineage	Rules, model/prompt versions, recommendation, confidence or threshold result.
Human involvement	Information presented, decision, modifications, comments and timestamp.
Execution	Tool, API, material parameters, target system, retries and response.
Business change	Document/object number, before/after values, posting identity and final status.
Outcome	Business result, exceptions, remediation, conformance and value measures.
Trace references	Correlation ID and immutable references to supporting records.

ILLUSTRATIVE AUDIT SUMMARY

TRANSACTION**AI-TXN-2026-009817**

A registered maintenance-risk agent, operating under policy EAM-RISK-2.1, detected a threshold exception for equipment 100045 and created notification 400018762 through an approved API. The agent identity, policy decision, data references, tool call, resulting document and monitoring outcome are linked by correlation ID CORR-EAM-74A19.

POLICY EAM-RISK-2.1	EQUIPMENT 100045	NOTIFICATION 400018762	CORRELATION ID CORR-EAM-74A19
--------------------------------------	-----------------------------------	---	--

Illustrative identifiers only.

PART 06

A business-process control model for agent transactions

Auditability begins before the agent is built. The process must be decomposed into tasks and decision points, and each task must be assigned an actor, authority boundary, control requirement and evidence obligation.

PROCESS-DESIGN QUESTION	REQUIRED DECISION
What business outcome is intended?	Define the measurable result and accountable process owner.
What tasks make up the process?	Identify reads, analyses, decisions, approvals, postings and exceptions.
Who performs each task?	Assign human, supervised agent, autonomous agent or system workflow.
Who has decision authority?	Separate recommendation, approval, execution and oversight.
What may the agent access?	Define objects, fields, plants, company codes, values and time boundaries.
What may the agent change?	Allow-list transactions, APIs, document types and permissible outcomes.
When must a person intervene?	Set thresholds for approval, uncertainty, exceptions and materiality.
What proves compliance?	Specify evidence, correlation, retention, reporting and review frequency.

Segregation of duties still applies

An agent should not collapse incompatible roles merely because it can technically perform multiple steps. The design must preserve distinctions such as request versus approval, recommendation versus posting, creation versus release, and configuration versus production oversight.

Authority should be task-specific

Agent access should be bounded by business purpose rather than copied from a broad human role. For example, an agent may display equipment and order history, calculate a risk score and create a draft notification, while release, budget commitment or safety-critical status changes remain human-controlled.

PART 07

Audit walkthrough: maintenance-risk intervention

The following example illustrates how process design and technical records work together in an SAP Enterprise Asset Management scenario.

STEP	ACTOR	ACTION	AUDIT EVIDENCE
1	Condition event	A threshold is exceeded for a monitored pump.	Event ID, equipment, timestamp and source.
2	Risk agent	Retrieves authorized history and open work.	Agent/session ID, data objects, tool calls and scope.
3	Risk agent	Evaluates policy and recommends priority.	Policy/model version, result and supporting factors.
4	Planner	Reviews and modifies the proposed scope.	User ID, information displayed, edits and approval.
5	Execution agent	Calls the approved notification API.	Agent identity, tool, parameters and API response.
6	Execution agent	Creates the maintenance notification.	Notification number, change record and posting identity.
7	Supervisor	Monitors completion and exception status.	Process events, service result and follow-up actions.

What the auditor should see

- The initiating event and the equipment record that caused the process to start.
- The agent's authorized scope and the information used in its evaluation.
- The exact recommendation presented to the planner.
- The planner's modification and approval, not merely a generic approval flag.
- The API call and resulting SAP notification linked through one correlation ID.
- Evidence that the transaction complied with the approved authority model.

AUDIT CONCLUSION

The evidence distinguishes what the event triggered, what the agent recommended, what the human decided, what the execution agent submitted and what SAP recorded.

PART 08

Controls across the agent lifecycle

LIFECYCLE STAGE	CONTROL OBJECTIVE	REPRESENTATIVE EVIDENCE
Discover	Select a legitimate, valuable business problem.	Use-case charter, owner and success measures.
Design	Define process, tasks, authority and human-agent boundaries.	Process model, RACI, control matrix and exception paths.
Build	Implement constrained tools, identities and data access.	Agent specification, API inventory and role design.
Test	Prove accuracy, security, control behavior and traceability.	Test scripts, negative tests, approvals and defect closure.
Approve	Accept residual risk and authorize production scope.	Business, security, architecture and risk sign-offs.
Operate	Monitor transactions, drift, exceptions and business value.	Logs, alerts, reviews, conformance and outcome metrics.
Change	Control prompt, model, tool, policy and scope changes.	Version history, regression tests and release approval.
Retire	Remove access and preserve required evidence.	Decommission approval, credential revocation and retention record.

Control testing should include failure paths

- The agent attempts a task outside its approved scope.
- The initiating user lacks authority that the agent possesses.
- A tool receives invalid, manipulated or excessive parameters.
- Required business context is missing or stale.
- A high-impact decision is attempted without the required human approval.
- Logs are incomplete or correlation identifiers do not reconcile.

PART 09

Implementation principles

01 Begin with the business process, not the model.

Define the outcome, actors, decisions and consequences before selecting agent technology.

02 Classify autonomy task by task.

A single process can contain informational, supervised and autonomous activities.

03 Use stable identities for humans, agents and tools.

Avoid shared technical users that obscure attribution.

04 Apply least privilege and purpose limitation.

Authority should match the specific task, data and business boundary.

05 Design evidence at the same time as the process.

Do not treat logging as a post-deployment activity.

06 Preserve human decisions as first-class records.

Capture what was presented, what changed and what was approved.

07 Correlate agent, integration and business records.

One trace reference should connect the entire transaction chain.

08 Monitor outcomes as well as uptime.

Auditability should show whether the process remained compliant and delivered value.

The target-state audit experience

A reviewer should be able to start with an SAP document, a person, an agent or a process instance and navigate to the same consolidated evidence. The interface should summarize the business event and link to authoritative source records rather than unnecessarily duplicating sensitive content.

DESIGN STANDARD

If the organization cannot explain a consequential agent transaction to a process owner, an auditor and a security reviewer using the same trace, the control design is incomplete.

PART 10

AlphaOak's role in creating audit-ready agentic processes

AlphaOak helps clients translate agentic AI capabilities into controlled business operations. Our responsibility is to make the business process, authority model, human-agent division of work and audit evidence explicit, then help the client configure, validate and operationalize those decisions in its SAP landscape.

10.1**Business-process and transaction design**

- Decompose end-to-end processes into tasks, decisions, handoffs, postings and exceptions.
- Identify where an agent informs, recommends, approves, executes, monitors or escalates.
- Define the SAP business objects, transactions and APIs involved at each step.
- Document expected outcomes, failure modes and operational dependencies.

10.2**Human and agent responsibility design**

- Classify each activity as human-only, supervised-agent, bounded-autonomous or system-controlled.
- Define accountable business owners, approvers, agent owners and technical operators.
- Develop RACI and decision-rights models that preserve segregation of duties.
- Establish escalation thresholds for risk, uncertainty, value and process exceptions.

10.3**Authority and policy definition**

- Translate business responsibilities into agent policies, role requirements and tool boundaries.
- Define permissible data, organizational scope, transaction types and write actions.
- Align agent authority with existing SAP security and business-control structures.
- Document autonomous operating envelopes and required human approvals.

10.4**Audit evidence and traceability requirements**

- Define the minimum evidence required for each material business task.
- Develop correlation, retention and reconstruction requirements across SAP and supporting services.
- Map evidence sources to business controls and audit questions.
- Design audit views and reports understandable to process, risk and technical stakeholders.

10.5**Control validation and audit readiness**

- Create positive, negative, boundary and exception test scenarios.
- Validate that human and agent permissions produce the intended effective authority.
- Confirm that recommendations, approvals, API calls and SAP documents reconcile.
- Facilitate control walkthroughs and evidence reviews with client stakeholders.
- Support remediation planning when policy, logging or process gaps are identified.

10.6**Governance enablement**

- Establish repeatable intake, classification, approval and change-control practices for agent use cases.
- Help define ownership forums spanning business, SAP, data, security, risk and audit.
- Develop operating procedures, role guidance and training for users and agent owners.
- Connect technical monitoring to process performance, compliance and realized value.

ALPHA OAK COMMITMENT

We help clients create a documented, testable chain from business responsibility to transaction authority, system execution and audit evidence.

Responsibility boundary

AlphaOak provides advisory, design, implementation and validation expertise within the agreed engagement scope. The client retains ownership of its business policies, regulatory interpretations, risk acceptance, production access approvals, control operation, data governance and final accountability for deployed agents and resulting business transactions.

CONCLUSION

Confidence comes from traceable responsibility

Agentic AI does not remove the need for established SAP controls. It makes the definition of actors, authority, decision rights and evidence more important. A trusted design must account for both human and agent participation, including autonomous activity executed under an approved policy.

The strongest audit model follows the business process from intent to outcome. It identifies who requested the work, which agent acted, what authority was applied, which information supported the decision, where human judgment entered, which tool executed the action and what SAP ultimately recorded.

FINAL PRINCIPLE

Every consequential agent transaction should be explainable as a business decision, attributable to human and agent identities, supported by approved authority, and reconstructable from durable evidence.

AlphaOak's contribution is to bring business-process expertise, SAP transaction understanding, role and authority design, human-agent operating models and audit traceability together. The objective is not documentation for its own sake. It is a controlled operating capability that protects the system, defines accountability and gives the client confidence that its AI-enabled processes can withstand operational, security and audit scrutiny.

A practical starting point

- 1 Select one material SAP process containing both human and agent activity.
- 2 Map the process at task and decision level.
- 3 Assign authority and evidence requirements to every step.
- 4 Trace one transaction end to end using current system records.
- 5 Close gaps before expanding agent autonomy or transaction scope.

REFERENCES

Sources

1. [SAP Architecture Center — Agentic AI & AI Agents on SAP Business AI Platform](#)
2. [SAP Architecture Center — AI Agent Behavior Mining with SAP Signavio](#)
3. [SAP Architecture Center — Third-Party MCP Server Access to SAP Solutions on BTP](#)
4. [SAP Architecture Center — SAP AI Agent Hub for Enterprise AI Governance](#)
5. SAP Help — Security Audit Log in SAP S/4HANA Cloud
6. [SAP Help — Read Access Logging](#)
7. SAP Help — SAP BTP Audit Log Service and Audit Log Retrieval API
8. SAP Help — Audit Logging in SAP Build Process Automation

Sources were accessed in September and October 2026. SAP product capabilities, availability and terminology change; implementation teams should confirm current product documentation and tenant-specific availability.

Important notice

This white paper presents a business and technology control framework. It is not legal, regulatory, accounting or independent audit advice. Control requirements should be reviewed with the client's legal, compliance, internal audit, cybersecurity and external assurance professionals as appropriate.



Trace one transaction before you widen agent autonomy.

AlphaOak helps utilities and asset-intensive organizations map one material SAP process, assign authority and evidence to every step, and close the gaps before agents go further.