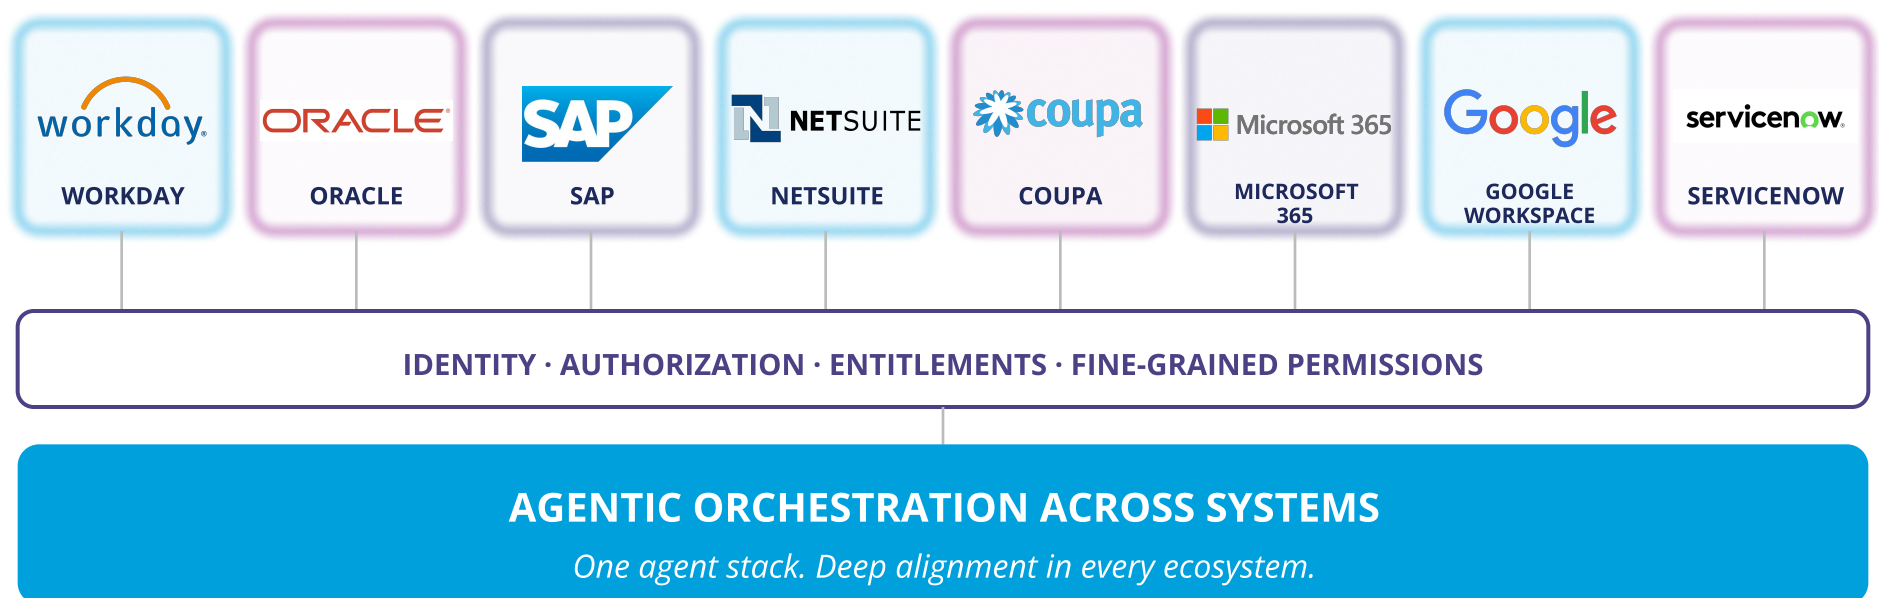


Multi-System Coverage and Governance Framework

Agentic work crosses your systems of record, so a vendor's credibility now comes from coverage, not pedigree inside any one platform. This is the framework for CFOs and CIOs evaluating that coverage, and the scorecard for choosing between a native agent and a specialized agentic solution.

YOUR SYSTEMS OF RECORD



Coverage is the new pedigree.

The three criteria for enterprise multi-system coverage.

1. Native interfaces, native object models = deep alignment.

No screen-scraping. No RPA. No additional databases. Instead, look for APIs, event streams, MCP servers, and agent gateways all working against each platform's actual object model.

2. A consistent agentic experience, system to system.

Validate that you will get the same agent behavior, override surface, and explainability trace regardless of where the user is logged in.

3. Orchestration across the systems, not within them.

Seamless connectivity is key. A vendor exception that begins in Coupa, resolves through Microsoft 365, posts to the Workday GL, and closes a ServiceNow ticket. All in one identity-bound flow.

When to trust a native agent. When to layer a specialized agentic solution.

Use this scorecard to evaluate your needs at the next vendor diligence meeting.

DIMENSION	A NATIVE AGENT IS ENOUGH	YOU NEED A SPECIALIZED AGENT
01 / SCOPE Where the work lives Are exceptions, vendor data, and approvals all inside one platform, or spanning many?	All inside one platform. Single ERP. Single P2P. One business unit, one geography, one stack.	Work spans two or more systems of record. Multiple ERPs, plus P2P, plus collaboration, plus ITSM.
02 / IDENTITY Identity propagation Whose identity does the agent carry into each system?	Same identity provider as the platform. Agent runs inside the platform's own SSO and entitlement model.	Identity has to propagate across multiple platforms. No global service accounts. No impersonation tokens.
03 / AUDIT Controller-grade log Can your CIO produce a single audit log of every agent action?	Platform's own audit log is sufficient. Auditor accepts the platform's existing compliance posture.	You need a unified log across systems. Who-did-what-where, queryable in one place by a controller.
04 / LEARNING Does it compound? Does day 365 look smarter than day 1 for your specific business?	Native agent's training and feedback loops are inside the platform. Sufficient for your scope.	Compounding requires the cross-system signal. Vendor patterns across ERPs that no single platform can see.
05 / TEMPO Mandate alignment Does the vendor have your finance-specific KPIs as a first priority?	Yes. The platform vendor's roadmap covers your KPI on your timeline, and you can wait for it.	No. Vendor mandate spans HCM to payroll to procurement. Finance specifics are a smaller share of focus.
06 / REFERENCE Production proof Who else runs this stack, at your scale, against your platform mix, today?	Native agent has many customers running it at your scale on the same platform, single-stack.	Specialized agent has multi-system production deployments at your scale, against your specific platform mix.

Ready to get trusted, multi-system coverage?

Let's talk. Email us at hello@auditoria.ai to schedule a call.

Or, visit us at auditoria.ai



Since 2019, Auditoria has set the standard for trusted, agentic AI in enterprise finance. Purpose-built for the Office of the CFO, Auditoria autonomously runs and optimizes your accelerating cash performance with the accuracy, control, and auditability your finance team demands. critical AP and AR workflows, accelerating cash performance with the accuracy, control, and auditability your finance team demands.