

GUIDE

Agentic AI: Due Diligence Guide for the C-Suite



Strategy Before Software

The 3- To 6-Month Decision Window: Are You Ready?

Agentic AI is moving from experimentation to enterprise-scale deployment, and the shift is happening faster than most HR technology transitions before it:

- Full **enterprise AI implementation has jumped 282%** year over year.
- **67% of CEOs now consider agent deployment critical** to staying competitive.
- Gartner projects that **40% of enterprise applications will include integrated, task-specific AI agents** by the end of 2026.

That means there's a narrow decision window for business leaders. You only have three to six months to define your agentic AI strategy before falling behind companies that move faster.

But that doesn't mean you need to rush into a software purchase. Gartner also predicts more than 40% of agentic AI projects will be canceled by the end of 2027 due to unclear value or weak governance. You need a defensible evaluation framework in place before the pressure to act forces a reactive buy.

Is the 81-Point Execution Gap Holding You Back?

The biggest risk to your AI deployment isn't the technology itself, but your data foundation. While 88% of organizations know orchestrating their people is critical, only 7% have made meaningful progress.

That 81-point execution gap is driven by traditional, isolated HR silos. If your recruiting, onboarding, and performance data live in disconnected systems, AI won't make you agile. It'll just make your fragmented processes fail faster.



What Does It Mean To Orchestrate Work?

Orchestrating work is the ability to organize and reorganize people, skills, data, and technologies around business-critical outcomes as your organization's needs change.

What This Guide Helps You Decide

This guide is designed to help you confidently decide where to deploy AI agents and how to choose the right vendor. Specifically, it will help you:

- **Define what actually qualifies as agentic**, and how to distinguish it from AI-assistant features rebranded as agents.
- **Identify which HR and people workflows** are mature enough for autonomous execution today.
- **Run vendors through a due diligence scorecard** that your CIO, CISO, and legal team will sign off on.
- **Quantify ROI, control operational and reputational risk**, and protect employee trust.
- **Prepare the workforce** — skills, roles, and the operating model — alongside the technology.

Start With Clarity: Understanding Assistants vs. Agents

Gartner's most-cited warning about the current market is blunt: the single most common misconception among buyers is labeling an AI assistant as an agent. They're calling it "agent washing" — the rebranding of chatbots, automations, and existing assistant features as agentic without the underlying autonomy to back it up. The firm estimates that only a fraction of vendors claiming agentic capabilities are actually delivering on them.

Practical definitions to use in vendor meetings

What it's called	What it actually does
AI Assistant	Simplifies discrete tasks (drafting, summarizing, searching) with continuous human input. The human remains both the doer and the decider at every step.
True Agentic AI	Operates autonomously across an end-to-end task: it can plan, execute, monitor outcomes, and iterate within defined constraints. The human shifts from doer to supervisor — handling exceptions, approvals, and governance rather than every click.

The decision-stage litmus test

Ask every vendor these four questions during your demos:

1. What task can your agent complete end-to-end without a human click path?
2. What decisions can it make on its own? What decisions are explicitly, technically prohibited?
3. How does it prove what it did? Ask to see the logs, audit trail, and rationale for a real action.
4. What is the escalation workflow when confidence is low, or the stakes are high?

If a vendor can't execute a complete workflow with guardrails live, you're likely testing an assistant feature wearing agent branding.



Where Agentic AI Creates Measurable Value in HR

Not every HR workflow is ready for autonomy, and treating them as equally ready is how pilots stall. Prioritize workflows that are high-volume and repeatable, rules-based with clearly defined policy constraints, and measurable in time, cost, risk reduction, or quality. The four categories below are where agentic AI vendors most often demonstrate credible, end-to-end autonomy today.

Category	What agents do today	ROI levers
Talent acquisition operations	• Candidate outreach sequencing and follow-up • Interview scheduling, reminders, and multi-party coordination • Screening workflow orchestration within compliant, defined criteria	• Time-to-fill • Recruiter capacity freed for high-touch roles • Candidate drop-off reduction
Employee service delivery (HR Ops)	• Tier-1 case resolution and ticket routing • Policy Q&A with built-in escalation and case creation for anything outside policy • Document generation and collection workflows (offer letters, verification requests, forms)	• Case deflection rate • Faster resolution time • Improved employee experience scores
Onboarding orchestration	• Equipment and system-access provisioning requests • Role-based checklist completion and automated nudges to managers and new hires • Manager enablement reminders and compliance-task completion tracking	• Time-to-productivity • Compliance adherence • First-90-days experience
Learning and internal mobility enablement (with governance)	• Personalized learning plans generated from role and skill data • Nudges tied to organizational capability priorities • Internal opportunity matching, with transparency controls so employees understand why a role was surfaced	• Retention • Internal fill rate • Skill velocity

Agentic AI Vendor Scorecard

How this works: Score each category 0–5 using the anchors below. Every category uses the same 0/1–2/3–4/5 scale, so a 3 means the same thing whether you’re scoring product capability or compliance. Score against one specific, named use case per vendor during your personalized demo, rather than in a general platform demo.

Before You Score: Capture Your Baseline

Before scoring any vendor, capture your baseline: cycle time, cost per transaction, error rate, satisfaction score, and compliance incident rate. Without a before, you can’t verify vendors’ ROI claims.

Universal Scoring Anchors

Score	What it means
0	Not present, or vendor couldn't demonstrate it live
1-2	Exists on paper or in roadmap, or was described but not shown working
3-4	Demonstrated live, works for the standard case, but has gaps (manual workaround needed, partial coverage, or unverified at your scale)
5	Demonstrated live, fully functional for your named use case, no material gaps

Category Rubrics

Category	Score	Anchor
Product Capability	0-1	Agent requires a human click at every step, or vendor showed slides instead of a live run
	2-3	Completes the task end-to-end live, but fumbles edge cases or needs a human to catch errors
	4-5	Completes the task end-to-end live, including an edge case you picked, with exportable outcome reporting

Category Rubrics

Category	Score	Anchor
Governance and Auditability	0-1	No audit log, or logs only readable by the vendor's own engineers
	2-3	Logs exist and are human-readable, but approval thresholds aren't configurable by you
	4-5	Full decision trail a non-technical reviewer can read, plus configurable approval workflows and real-time alerting
Risk and Compliance	0-1	Bias testing is out of scope or coming soon; no data residency options
	2-3	A named bias-testing method exists but no reviewable report; standard security certs in place
	4-5	Reviewable bias-testing report, data residency matches your requirements, compliance posture covers every jurisdiction you operate in
Integration and Time-To-Value	0-1	Custom-build-only, no credible timeline
	2-3	Prebuilt connectors exist but implementation is 16+ weeks or requires IT to build config, not HR
	4-5	Prebuilt connectors to your core systems, under 12–16 weeks, HR and IT can jointly configure
Adoption and Experience	0-1	No change management materials beyond a training video; no AI-literacy resources; no references available
	2-3	Change management toolkit exists, but references are limited or not at your scale; no reskilling materials
	4-5	Full change management toolkit, verifiable adoption metrics, references at your scale, plus AI-literacy or role-transition materials you can adapt internally

Your Vendor Scorecard

Category	Weight	Vendor A	Vendor B	Vendor C
Product Capability	x1			
Governance and Auditability	x1			
Risk and Compliance	x1			
Integration and Time-To-Value	x1			
Adoption and Experience	x1			
Total (/25)				

Decision rule: Only advance vendors that score 18/25 or higher. Avoid advancing those who score below 3 in Governance or Risk and Compliance specifically. A strong product score can't offset a governance or compliance failure, even if the total clears 18.

Internal Stakeholders: Who To Involve Before You Sign

Agentic AI fails when treated as an isolated software purchase instead of a team decision. Bring in key stakeholders during discovery rather than after the contract is drafted.

Stakeholder	What they own in this decision
CHRO / People Leadership	The value case, the trust framework, and the operating-model changes required to make the deployment stick.
CIO / IT	Architecture, integration feasibility, and security controls across existing systems.
CISO / Security	Vendor risk assessment, access controls, monitoring, and incident response readiness.
Legal / Compliance	Privacy obligations, labor regulation exposure, and auditability requirements.
Finance	Independent validation of ROI assumptions and the value-capture model.
Employee Comms / Change Management	Transparency with employees, adoption planning, and protecting the experience during rollout.

Next Step: A Structured Software Evaluation

Don't jump into another open-ended pilot. Use this phased approach to guide your team through a thorough evaluation:

- 1. Assess and Align:** Figure out exactly where your data lives and where it breaks down before you talk to vendors. Find one or two workflows with the clearest ROI to anchor your first deployment.
- 2. Set Your Non-Negotiables:** Decide on your strict AI rules right now, like where a human must approve an action and what audit logs you need to see. Use these as a pass/fail filter.
- 3. Test the Integration:** Make shortlisted vendors prove that a single data point (like source of hire) actually flows end-to-end into onboarding and performance. Make sure their AI isn't just a layer slapped on top of disconnected tools.
- 4. Prep Your People:** Don't treat this like a standard software swap. Build your change management plan before you sign a contract, not after.

Turn HR Complexity Into Clarity

Fragmented HR tech forces AI to guess at answers using disconnected data. But when you treat your organization as a single unified talent platform, your agents become incredibly powerful.

ClearCo connects hiring, onboarding, performance, and learning in one secure place. Because we use a shared data foundation, built-in capabilities like the Sourcing, Interview, and Review Agents can instantly act on clean data. You gain end-to-end autonomy without the messy integration debt.



Stop settling for fragmented HR tools.

[Book a ClearCo demo](#)



ClearCo

clearcompany.com