From Exposure to Value: A New Diagnostic for the Agentic Era

sia-partners.com

July 2026

/ Table of contents

When Sia published the first AI OrgReview in September 2024, the defining question was: *which tasks could a language model accelerate?* That was the right question for its time. Organizations were just beginning to understand what generative AI was capable of, and exposure rubrics – borrowed from OpenAI and University of Pennsylvania research – provided a credible first lens.

Eighteen months later, that question has been answered. The models have taken a major step forward. In early 2026, every major technology actor – Anthropic, OpenAI, Google DeepMind, Meta – is competing on frontier capabilities, and the pace of release has moved from annual to monthly. Meta's April 2026 model launch signals that even actors previously absent from the frontier race are now competing directly. The pace of change will never feel this slow again.

The question organizations face today is not *can AI impact this task*? The answer is almost always yes. The question is: **how do we capture the value – on the P&L, in measurable business outcomes, before our competitors do?** Capitalizing on this opportunity will require investing in skills, redesigning workflows, and supporting transformation as employees adapt to working alongside agents and intelligent machines.

This update to AI OrgReview reflects that shift – deepening the task-level exposure analysis that has always been at its core, and extending it toward workflow design, business outcome quantification, and investment-grade diagnostics. We are, as one of our own leaders put it, at kilometre one of a marathon. This paper is designed to help organizations run it.

1. What Has Changed Since September 2024

1.1 The models crossed a threshold

The 2024 version of AI OrgReview was grounded in the capabilities of a specific generation of models. That generation has been superseded – not iteratively, but in a step-change. As of early 2026:

- **Autonomous agent execution has moved from research demos to enterprise production.** Cox Automotive deployed AI agents at scale on Amazon Bedrock AgentCore to automate vehicle inventory and pricing workflows across its dealer network. AT&T rebuilt its customer care operations around NVIDIA NeMo-based agents handling multi-step service requests autonomously. Blue Origin used AI agents to design hardware components for lunar missions – compressing months of engineering iteration into days. These are not pilots; they are operating at scale.

- **Reasoning capabilities have moved tasks previously requiring human judgment firmly within model reach.** What changed is not just raw performance but reliability under ambiguity: models can now operate on incomplete instructions, flag inconsistencies, and escalate appropriately – the minimum bar for deployment in regulated environments.
- **The cost of running frontier AI has fallen by 67–80% in under two years.** In 2023, processing one million tokens with a frontier model cost upwards of $30 (OpenAI GPT-4 at launch). By early 2026, equivalent or superior capability costs $3–5 per million tokens. Sub-frontier models capable of handling high-volume, structured tasks now cost less than $0.50 per million tokens. This makes continuous, workflow-embedded AI economically viable in a way that was not true eighteen months ago.

The 2020 "scaling law" findings continue to hold: model performance improves with compute investment, and the investment has not slowed. Organizations should assume continued capability increases through the planning horizon of any transformation they undertake today.

1.2 The enterprise reality: a widening divide

The data tells a paradoxical story. Adoption is broad – but impact remains shallow. 71% of large organizations are using generative AI in at least one function, and less than 50% of employees report regular use at work, up from 20% in 2023. Yet more than 80% of organizations report no tangible EBIT impact from their AI investments to date, and only 21% have fundamentally redesigned any of their workflows. On productivity, empirical measurement puts the average gain from AI assistance at +15%, rising to +26% on software development tasks.

The gap between the 71% who have *deployed* and the 21% who have *redesigned* is precisely where value is lost. Tool adoption does not create business impact. Changing how work gets done does – and that is a different kind of problem entirely.

1.3 From "how much time do we save?" to "how much value do we create?"

We are entering a second phase of enterprise AI adoption. The first – roughly 2023 to 2025 – was about productivity: doing existing work faster with AI assistance. The second, which leading organizations are already navigating, is about *process reinvention*: redesigning how value is created, not just how quickly existing work gets done[0]. The analogy holds: the internet's first phase gave us faster letters. The transformative phase – e-commerce, platform business models – required reinvention from the ground up, arrived later, and delivered disproportionately more value.

The workforce stakes of this shift are substantial. Our conviction, grounded in the data we observe, is that the challenge facing organizations is not primarily a technology challenge – it is a redesign challenge. Millions of roles will require significant transformation every year through 2031. This is not a jobs apocalypse; AI will create more jobs than it eliminates over that horizon. But continuous role redesign at that scale is the 20x larger organizational challenge, and most organizations are not positioned to manage it.

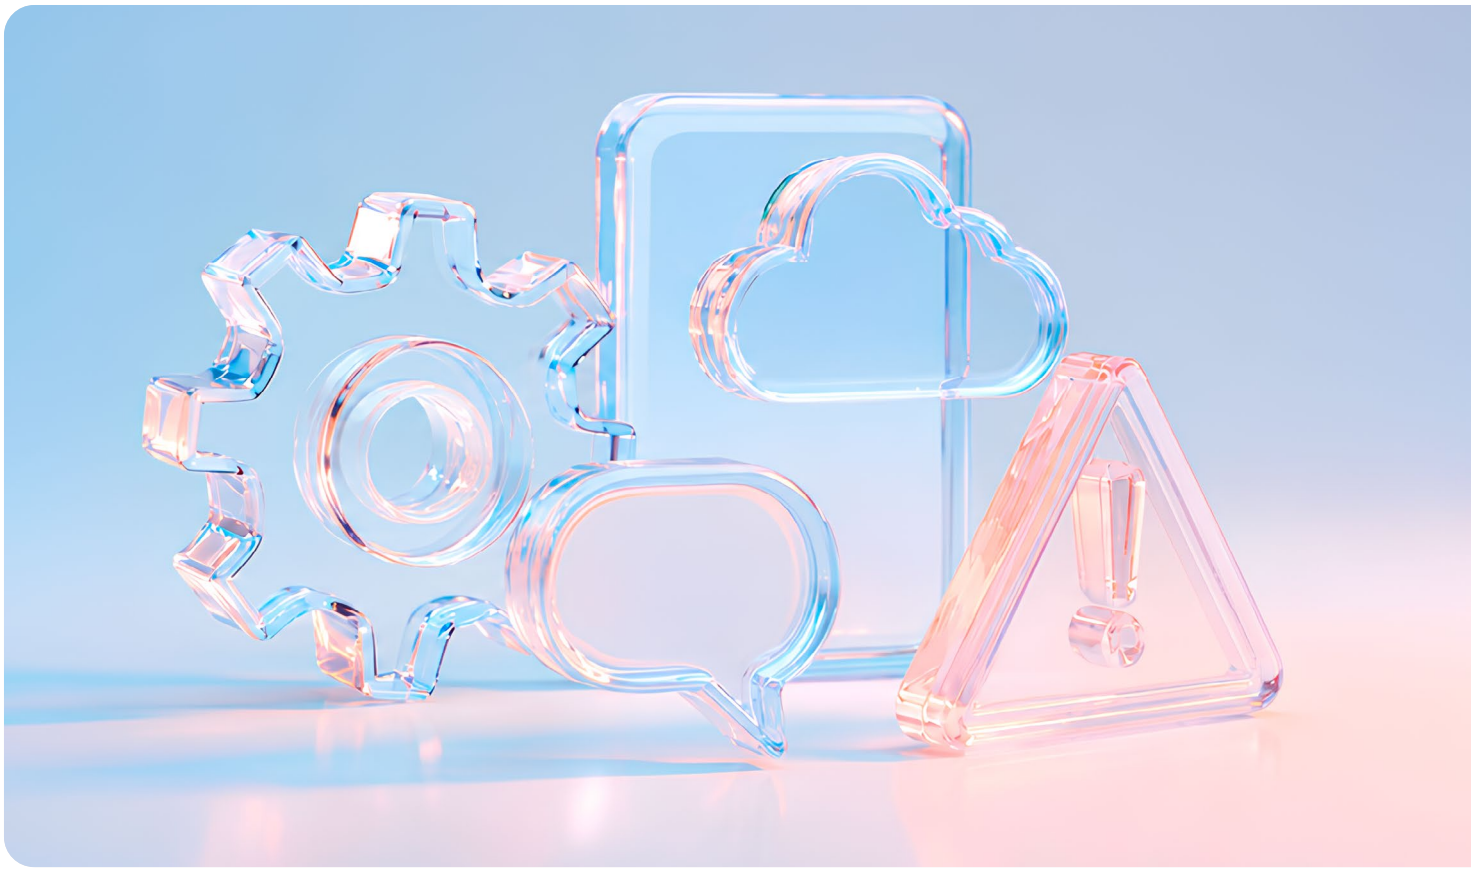

2. Updated Methodology

2.1 From the Exposure Rubric to the Value Impact Model

The original AI OrgReview used a three-level rubric (No Exposure / Direct Exposure / Advanced Exposure), calibrated around a single question: could a task be completed at least 50% faster using a language model?

This rubric remains a useful first filter, but it is insufficient for three reasons. First, it does not distinguish between augmentation – where a human performs a task faster with AI assistance – and automation – where an AI agent executes the task with humans in an exception-handling role only. These have radically different implications for FTE planning, cost modeling, and operating model design. Second, it does not capture agentic impacts, where entire workflows – not individual tasks – are restructured around AI execution. Third, it does not connect exposure to financial value, which is what organizations and their leadership need to make decisions.

Our updated framework introduces four impact levels designed to produce actionable financial estimates:

Level	Definition	Workforce Implication	Value Formula	Illustrative Example
E0 Uniquely Human	Task performed entirely by a human with no AI involvement. The human is fully accountable for execution and output. No AI assistance is used.	No change to FTE count or role design; establishes the cost baseline for comparison with AI-enabled levels.	N/A – establishes the cost baseline against which AI-enabled savings are measured.	Strategic advisory, stakeholder negotiations, ethical judgment calls, or crisis management where human judgment, empathy, and accountability are non-substitutable.
E1 AI-Automated	Task previously executed by a human is now delegated to an AI agent, with humans in a review or exception-handling role only. The agent initiates, executes, and delivers the output.	Structural FTE reduction or reallocation possible; requires role redesign and a governance framework.	(Tasks per FTE × automation rate × FTE cost) – (governance + monitoring overhead) = net value.	A customer service workflow where 60–70% of interactions are handled autonomously, as observed in leading financial services deployments.
E2 AI-Augmented	Task performed by a human, significantly faster or at higher quality, with AI assistance. Human remains accountable for the output. Time savings ≥50%.	Same FTE count; productivity increase; potential for scope expansion or redeployment.	Productivity gain × loaded FTE cost × impacted headcount = value pool.	A compliance analyst using AI to compare regulatory frameworks across jurisdictions. Work previously requiring 3 months and 8 consultants (Sia's UBS engagement) now takes 6–8 weeks with a smaller team.
E3 AI-Redesigned	The process or workflow itself has been fundamentally restructured around AI execution. The prior human-designed workflow is obsolete. New roles emerge: agent supervisors, workflow architects, decision validators.	New operating model required; existing roles partially or wholly redefined; governance is a first-class deliverable.	Revenue impact + cost transformation + speed-to-market advantage; modeled at business unit or P&L level.	Trade lifecycle compression in capital markets – moving to same-day settlement requires AI agents to handle operational exceptions in near-real-time, replacing workflows that previously required 24-hour human analyst cycles.

2.2 Enriched Segmentation

Beyond occupation-level analysis, the updated AI OrgReview segments impact across three dimensions:

- **/ By task type: Cognitive/knowledge work carries the highest E2–E1 impact. Relational and judgment-intensive work has lower automation potential but high augmentation value. Physical and manual work is context-dependent, with emerging impact via computer vision and robotics.**
- **/ By organizational layer:** The "sandwich" dynamic observed across our client base – where AI-native junior profiles and senior domain experts compress middle management – requires explicit modeling. Impact on individual contributor, managerial, and executive layers differs materially.
- **/ By sector-specific constraints: Regulated industries face hard limits on autonomous AI in high-stakes decisions. EU AI Act risk classifications and GDPR constraints must be incorporated into any serious financial model for E1 and E3 scenarios.**

3. What's Actually Holding Organizations Back

Finding 1: Exposure has increased – but the binding constraint has shifted

Applying the updated E2/L2/E3 methodology to an expanded occupation sample (850+ activities from O*NET 2025, validated through 40+ SME interviews and client engagement data, calibrated against model capability benchmarks as of Q1 2026), our analysis revises upward the exposure estimates from the 2024 paper:

- **E2 (AI-Augmented): 52–65% of reviewed activities can be completed significantly faster with current AI assistance – up from 36–46% in 2024**
- **E1 (AI-Automated): 18–28% of activities are candidates for autonomous agent execution with current or near-term technology**
- **E3 (AI-Redesigned): 8–15% of organizational workflows can be fundamentally restructured around AI within a 12–24 month horizon, contingent on operating model readiness**

These numbers matter less than what they reveal about where the real challenge lies. In 2024, the question was whether AI could handle a task. In 2026, it almost always can. The binding constraint has shifted to value capture – and across our 40+ SME interviews and client engagements, the pattern is consistent.

> We are not stuck on the technology. We are stuck on the organization, the data ownership, and the business leader who needs to own the outcome – not IT."

David Martineau
Sia Chief AI Officer, April 2026

Organizations that close the gap between deployment and P&L impact share three identifiable characteristics.

Governance at the right level: the transformation agenda is owned by a business leader, not housed in IT. When AI leadership reports into strategy or the COO function, the conversation becomes about outcomes; when it reports into IT, it stays technical.

Business-first framing: initiatives defined by outcome ("we're cutting regulatory review time from six months to six weeks") consistently outperform those defined by technology ("we're deploying an agentic workflow for compliance").

A value funnel, not a use case catalogue: 10–20 prioritized value pools with explicit KPI contracts and named business owners, rather than a broader catalogue of lower-conviction experiments with no hard qualification gates.

Finding 2: The half-life of workforce skills is collapsing

We see a structural shift that renders static workforce planning obsolete: **the half-life of technical skills will shorten to 2–5 years by 2030, down from 8–12 years today**. The World Economic Forum estimates 39% of workers' core skills will need to change by 2030.

The implications are direct. A workforce diagnostic conducted today will need to be refreshed within 18–24 months. Static AI impact assessments that produce a one-time FTE reduction estimate misrepresent the actual challenge, which is *continuous role redesign*. The primary capability organizations need to build is not "AI literacy" – that is a training solution to a structural problem – but *adaptive capacity*: the organizational muscle to continuously redesign roles and workflows as AI capabilities evolve.

The risk flagged consistently by our sector experts is *homogenization by* the mean: organizations relying on AI to make average employees perform at a higher average level, without investing in the senior judgment and domain expertise that AI cannot replicate. "AI will not substitute a doctor, but a doctor who uses AI will substitute one who doesn't" – the same dynamic applies across professional services, banking, and beyond.

Finding 3: Autonomous agents are entering production – governance is the new frontier

The agentic paradigm is no longer theoretical. Organizations that have successfully moved from GenAI augmentation to agentic deployment share a common characteristic: they invested in governance infrastructure before scaling.

The most underestimated risk identified in our sector interviews is **technical change management in an agentic environment**. Unlike static software, agents require continuous maintenance as business processes evolve. Deploying an agent and failing to maintain alignment between its behavior and evolving business logic creates liability, not value.

Three governance requirements are non-negotiable for L2/E3 deployments. Identity and access management: an agent acting on behalf of a user must be constrained to that user's actual permissions, not the administrator's. Explainability at decision points: in regulated industries, multi-step agentic workflows must demonstrate the same reliability standards as human judgment – and they are harder to audit than single-step outputs. Kill switches and rollback protocols: the ability to deactivate or revert an agent's scope without disrupting the broader workflow is an architectural requirement, not an afterthought.

Finding 4: Redesigning workflows is where AI value is captured

While much of the discussion around AI focuses on individual tasks, the real opportunity lies in redesigning how work gets done end to end. This was one of the original insights behind OrgReview: understanding which activities within a role can be accelerated or automated by AI provides a powerful starting point for transformation. However, task level analysis alone is rarely sufficient to capture the full value.

Most organizational workflows were designed for a world in which people completed every step manually. Applying AI to isolated activities within those workflows often delivers only incremental benefits because the broader process remains unchanged. Organizations may reduce the effort required to complete individual tasks, while still maintaining the same handoffs, approval layers, coordination requirements, and organizational boundaries that existed before AI.

This helps explain why many organizations have successfully deployed AI tools, yet relatively few have realized material business impact. Improvements at the task level do not automatically translate into meaningful improvements in cost, speed, quality or customer experience. The largest gains are often unlocked not through automation itself, but through redesigning the workflows that automation enables.

Capturing this value requires organizations to rethink how work flows across teams, roles, and systems. Activities that were historically distributed across multiple functions can be consolidated. Handoffs can be eliminated. Decision-making can be pushed closer to the point of action. The entire process steps can be simplified or removed. AI becomes exponentially more powerful when it is embedded within a redesigned workflow rather than applied to individual tasks.

As workflows change, so do roles. Work is redistributed between people and technology, reshaping responsibilities, required capabilities and where human expertise creates the greatest value. In many cases, the most significant impact of AI is not that jobs disappear, but that jobs evolve. This creates a challenge for leadership teams. Identifying AI use cases is relatively straightforward. Understanding their cumulative impact on workflows, roles, teams, and organizational structures is significantly more complex. Organizations need a way to move beyond isolated experiments and develop a fact-based view of how work itself should be redesigned.

This is where OrgReview provides a critical foundation. By analyzing the activities performed within roles and assessing their exposure to AI, organizations can identify not only opportunities for automation but also for redesigning workflows, reshaping roles, and rethinking operating models. Rather than treating AI adoption as a technology initiative, OrgReview helps organizations understand its implications for work, workforce, and organizational performance, creating a practical roadmap for AI-enabled transformation.

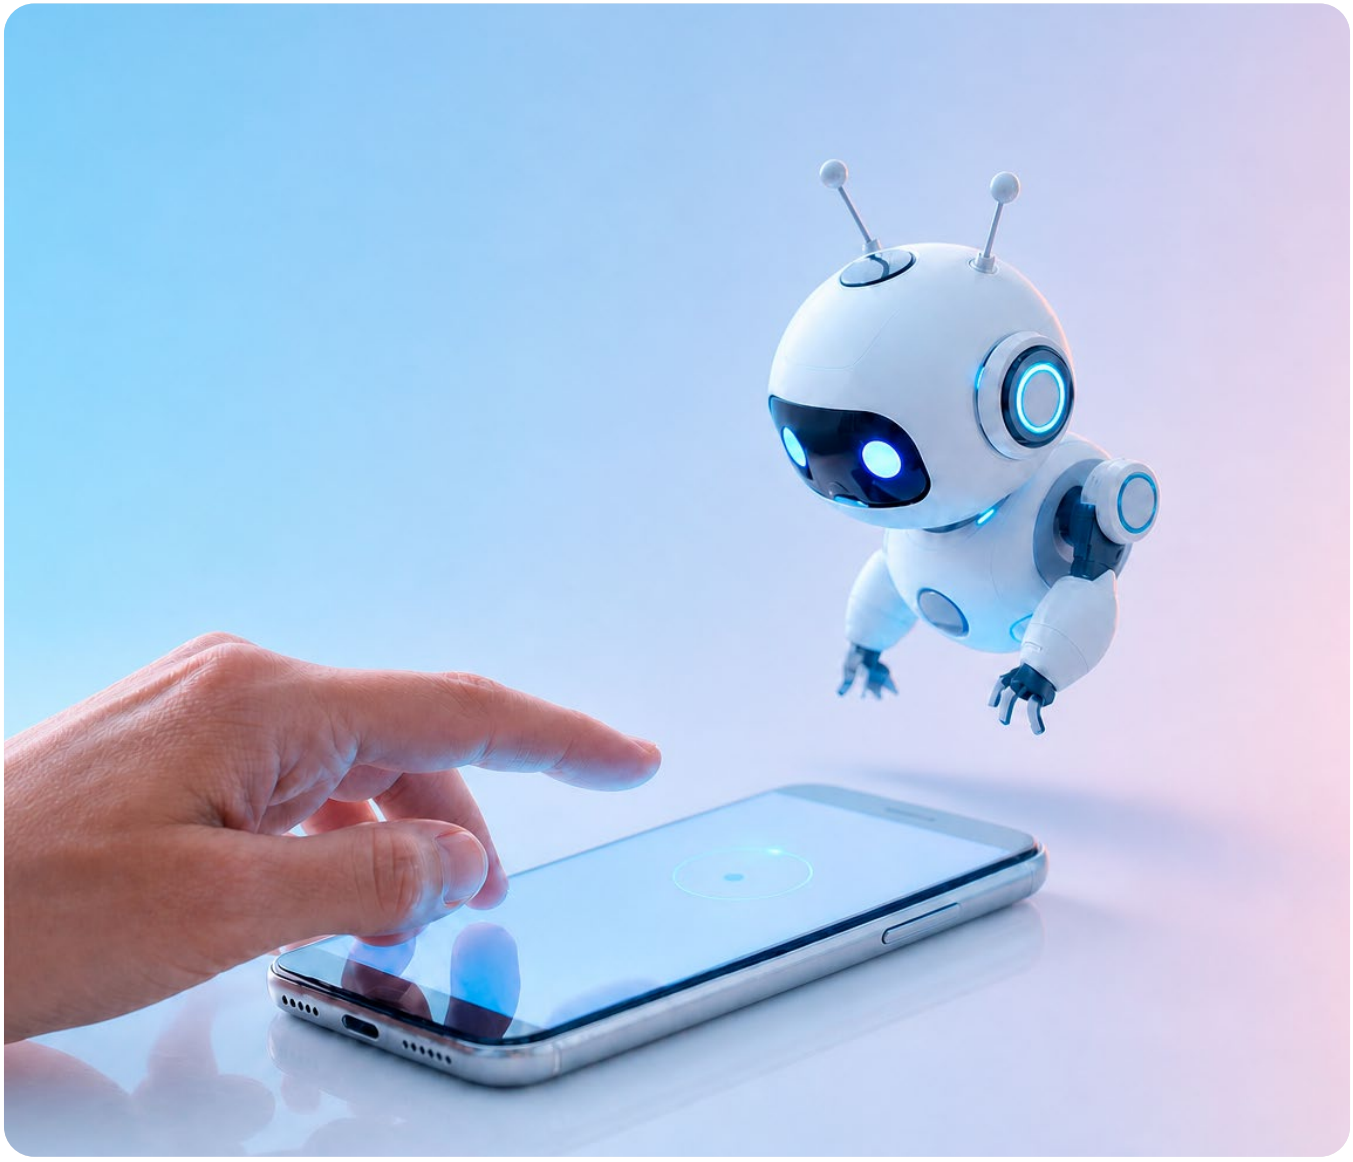

4. What Leading Organizations Are Achieving

For each function, we report observed impact ranges at organizations operating at leading practice – defined as those with strong data infrastructure, clear process ownership, and active adoption programs in place. These benchmarks are drawn from Sia's direct client engagement data across 800+ production AI deployments, validated against expert interviews and cross-referenced with published functional AI research.

A 30–50% discount should be applied to E1 and E3 estimates for organizations that have not yet resolved data readiness and operating model ownership.

Function	Primary AI value levers	E2 Augmentation (observed range)	E1 Automation potential	Leading-practice impact	Key constraint
Finance & Accounting	Reporting automation, anomaly detection, FP&A narrative generation, AP/ AR processing	60–70% of tasks	25–35%	20–35% reduction in time spent on transactional and reporting workflows; 15–25% FTE reallocation toward higher-value analysis	Data quality and ERP integration maturity
Legal & Compliance	Regulatory gap analysis, contract review, policy mapping, audit trail generation	65–80% of tasks	20–30%	Regulatory review time reduced by 60–80% in document-intensive workflows; Sia RegAI: UBS regulatory gap analysis from 3 months/8 consultants to 6–8 weeks with smaller team	Human validation required for regulatory decisions; high-risk classification under EU AI Act
HR Operations	Tier-1 query handling, onboarding automation, CV screening support, learning content generation	55–70% of tasks	30–45%	40–60% of tier-1 HR queries handled autonomously in mature deployments; significant productivity gains in learning content design	EU AI Act high-risk classification for hiring tools; strong data sensitivity constraints
Software Engineering	Code generation, automated testing, documentation, code review, security scanning	55–70% of tasks	20–30%	+20–40% task completion rate; in the right configuration, productivity velocity boosts "orders of magnitude greater than traditional development"	Cost of token usage at scale; silent regressions in agentic workflows require test-driven governance
Customer Service	Automated query resolution, case routing, real-time agent assist, voice automation	50–65% of tasks	35–55%	60–70% of standard interactions handled autonomously in leading financial services deployments; significant reduction in average handling time	Voice automation at enterprise scale technically complex; tolerance for error near-zero in sensitive interactions

Function	Primary AI value levers	E2 Augmentation (observed range)	E1 Automation potential	Leading-practice impact	Key constraint
Banking Operations	Trade settlement, fraud detection, KYC/AML, investment research, mortgage processing	45–60% of tasks	30–45%	41% of banking CIOs have deployed or plan to deploy AI agents by 2026; leading deployments: BNY multiagent platform, J.P. Morgan Legal Agentic Workflows (LAW), Mastercard Agent Pay	Regulatory explainability requirements; zero tolerance for data fabrication in financial reporting
Insurance	Claims processing, underwriting support, fraud detection, policy document handling	50–65% of tasks	35–50%	16% of insurers have deployed AI agents; 35% plan to by end 2026; simple claims (e.g., glass breakage) now automatable end-to-end; material reduction in call-center handling time via pre-call data collection agents	Complex/bodily injury claims must remain human; regulatory constraints on automated underwriting decisions
Procurement & Supply Chain	Supplier analysis, contract review, demand forecasting, exception management	45–60% of tasks	25–40%	Significant gains in document-intensive processes (RFP analysis, contract comparison); agentic scheduling and exception handling emerging	Highly context-dependent; exception logic often undocumented and held in human knowledge
Marketing & Content	Copy generation, creative variation, personalization at scale, market research synthesis	60–75% of tasks	20–35%	Cost of standard content production approaching zero; value shifts to personalization architecture and data-driven segmentation	Exceptional creative talent still commands premium; brand consistency governance required for agentic content

5. What This Diagnostic Does Not Capture

Model capability assumptions: Our E2/L2/E3 framework reflects capabilities as of early 2026. Given the pace of model development, E1 and E3 exposure estimates are likely to be conservative within 12–18 months. We recommend treating this diagnostic as a living document, refreshed annually.

Exposure ranges are Sia's own analysis, not a published dataset: The E2/L2/E3 percentages in Finding 1 are derived from Sia's application of the updated methodology to an expanded occupation sample, not a single third-party dataset. They represent observed impact ranges validated through client work and expert interviews, and should be read as directional estimates, not statistical averages.

Organizational readiness is not captured by the task model: Our analysis identifies the value *potential* of AI deployment. It cannot predict capture. Organizational factors – leadership commitment, data infrastructure quality, change capability, regulatory environment – are assessed separately and must be weighted explicitly in any financial model.

Sector heterogeneity: A financial institution under ECB supervision faces materially different constraints on L2/E3 deployment than an FMCG company. EU AI Act risk classifications must be incorporated into any compliance-sensitive model.

The "AI co-worker" effect is not fully modeled: AI agents as nominal team members – with persistent identity, memory, and task ownership – create organizational design questions our framework addresses at the E3 level but does not fully resolve. Role definitions, reporting structures, and accountability frameworks for human-AI teams remain a rapidly evolving area.

Data infrastructure dependency: At all levels of the framework, data quality is a binding constraint. A company without clean, owned, governed data cannot realize the value our model suggests.

6. What Leading Organizations Do Differently

Commission the diagnostic as a strategic input, not an HR exercise.

AI OrgReview is most valuable when its output feeds directly into the strategic plan and operating model review – not when it runs as a standalone workforce initiative. The question it answers – where is value concentrated, and what is the realistic financial case for capturing it? – is a board-level question. Treat it as one.

Use AI as a catalyst for redesigning work, not simply automating tasks

Many organizations are still approaching AI as a collection of disconnected use cases. While this can generate incremental productivity improvements, it rarely delivers transformational value. The greatest opportunities emerge when organizations redesign end to end workflows, removing unnecessary handoffs, simplifying decision making and fundamentally changing how work is performed. AI should be viewed as an enabler of work redesign rather than a tool for automating individual activities.

Prioritize workflow redesign

Most workforce transformation efforts begin by asking how jobs will change. In practice, organizations should start one level higher by understanding how core workflows will evolve. Once future workflows are defined, the implications for roles, skills, structures and headcount become significantly clearer. Attempting to redesign roles without first redesigning workflows often results in organizations replicating legacy ways of working with new technology.

Treat workforce transformation as an operating model challenge

The impact of AI extends beyond productivity. As work is redistributed between people and technology, organizations will need to rethink responsibilities, decision rights, governance, team structures and capability requirements. These are operating model questions, not technology questions. Leadership teams should ensure AI transformation is connected directly to broader discussions about organizational effectiveness and future ways of working.

Design for value capture from day one, not after the pilot.

The organizations that struggle have a common pattern: they build the use case first and design for value later. The sequence that works inverts this. Define the business outcome and its KPI. Identify the AI-enabled process that moves it. Design the organizational changes required for adoption. Then build. In that order.

Make governance an architectural decision, not an afterthought.

As deployments move from E2 to E1 and L3, governance complexity scales non-linearly. Access controls, explainability, rollback protocols, and continuous alignment between agent behavior and evolving business logic are not compliance tasks – they are engineering and operating model requirements. Organizations that wire them in from the start scale faster and with fewer incidents than those that retrofit them.

Invest in adaptive capability, not one-time training.

Given that the half-life of technical skills will be 2–5 years by 2030, the right response to workforce transformation is not a training program – it is a structural investment in organizational adaptability. The benchmark question is not "have our employees completed an AI literacy course?" but "can our organization redesign roles and workflows faster than our competitors?"

Plan for continuous refresh.

The organizations that create lasting advantage from AI will not be those that deploy the most tools. They will be those who become better at continuously adapting how work is organized and delivered. Workflow redesign, role redesign, and capability evolution should become ongoing organizational disciplines rather than one-time transformation activities. The ability to redesign work faster and more efficiently than competitors may become one of the defining capabilities of the AI era. Any AI workforce diagnostic conducted today will need to be revisited within 18–24 months. Build that cadence into the program from the outset. The organizations ahead in 2028 will not be those that did the best diagnostics in 2026 – they will be those that made it a continuous organizational practice.

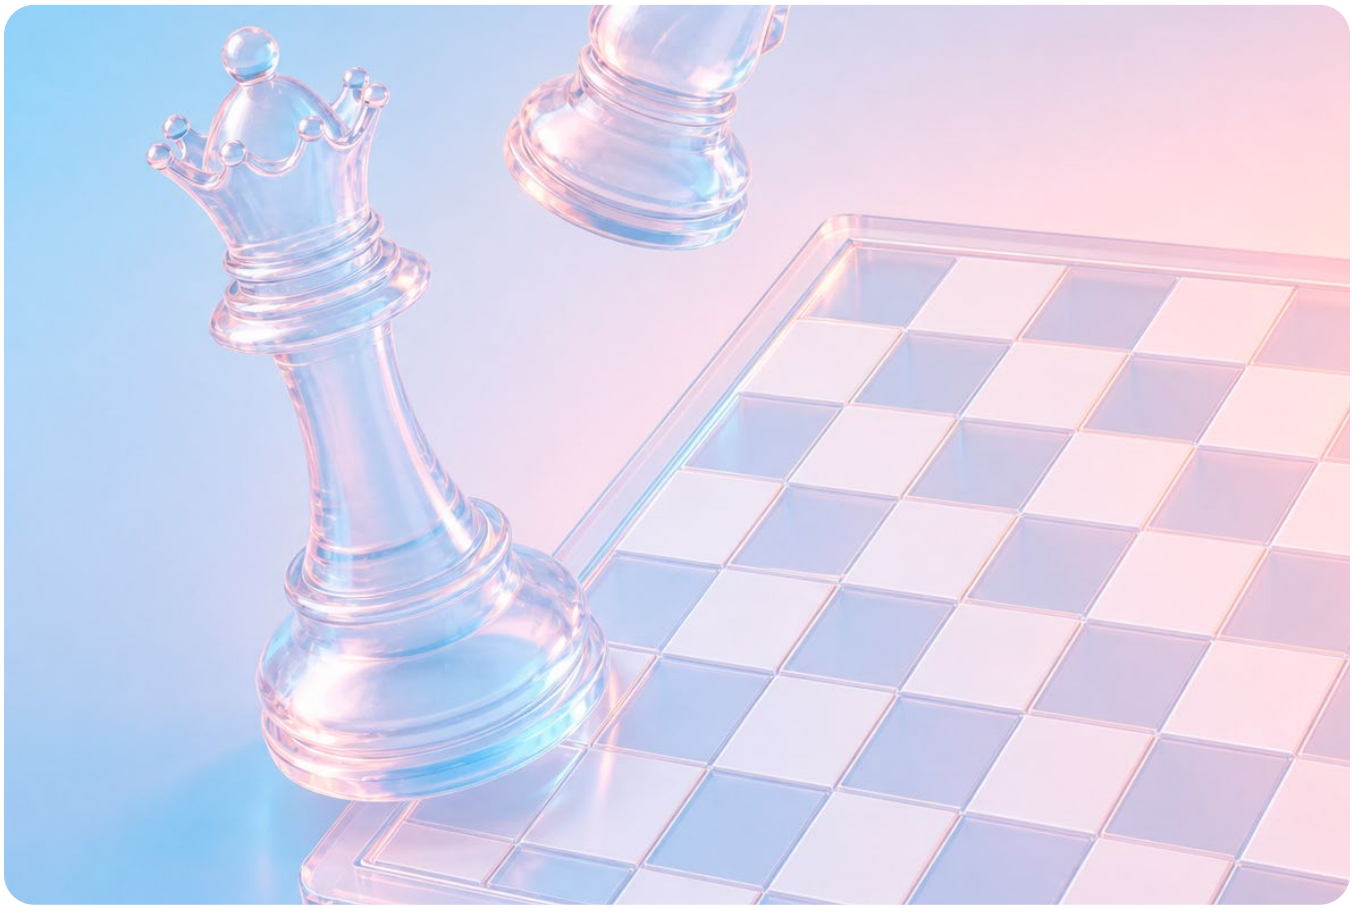

The first AI OrgReview answered a question that mattered in 2024: how exposed is this workforce to the capabilities of generative AI? That diagnostic framing served its purpose. The 2026 context demands a different one. The models have advanced. The divide between organizations that deploy AI and those that capture value from it is widening. And the emergence of autonomous agents has made the unit of analysis the workflow – not the task – and the measure of success the P&L impact, not the exposure score.

The organizations that create the greatest value from AI will not necessarily be those with the most advanced models or the largest technology budgets. They will be those who are able to redesign key workflows, reshape roles, and adapt their operating models faster than their competitors.

This requires a different level of analysis. Leaders need to understand not only which activities can be automated, but how workflows will change, where human judgment creates value, and what those changes mean for roles, teams, and organizational structures. In many organizations, this represents one of the largest transformation opportunities of the next decade.

AI OrgReview v2 was built for this challenge. Whether you are a CHRO designing a workforce transition program, a COO building an agentic operating model, or an executive team assessing where AI creates the most value in your business, this diagnostic provides the foundation for decisions that are evidence-based, financially grounded, and built to hold up as the technology continues to evolve.

Ultimately, the question is no longer whether AI will change the way organizations operate. The question is which organizations will be able to redesign work fast enough to take advantage of it.

Updated Scoring Framework: Illustrative Examples

Occupation	Activity	Level	Rationale
Financial Examiner	Review balance sheets to confirm institution assets	E2	AI accelerates data parsing and pattern identification; human judgment required for final assessment
Compliance Manager	Map regulatory differences across frameworks	E2–E1	AI automates initial comparison; human validation required for regulatory decisions
HR Specialist	Process standard onboarding queries	E1	Autonomous agent handles ≥60% of tier-1 HR queries in leading deployments
Legal Secretary	Schedule and make appointments	E1	Calendar agents fully automate this in current deployments
Business Intelligence Analyst	Generate standard reporting dashboards	E2–E1	AI augments significantly; autonomous generation emerging for standardized reports
Customer Service Representative	Handle standard product/ service queries	E1–E3	60–70% autonomous handling observed; full redesign possible with integrated backend systems
Software Engineer	Code generation and review	E2–E1	+20–40% task completion; agentic coding shifting toward E1 for defined scopes
Tax Preparer	Prepare standard tax filings	E1	High automation potential for standard cases; human judgment required for complex or advisory decisions
Travel Agent	End-to-end trip planning and booking	E1–E3	Agentic commerce paradigm applies directly; machine execution with human confirmation at key decision points

/ Sources

/ Panorama 2026 de l'IA en Entreprise*, French Tech Grand Paris / VivaTech, 2026.

/ Gartner, *Predicts 2026: AI Workforce Impact*, November 2025. The same report projects the half-life of technical skills shortening to 2–5 years by 2030, down from 8–12 years today.

/ Sia's own analysis applying the updated E2/L2/E3 methodology to 850+ occupational activities sourced from O*NET 2025, calibrated against model capability benchmarks as of Q1 2026, and validated through 40+ SME interviews and client engagement data (2025–2026). These figures represent observed impact ranges, not a statistically representative sample. The 2024 comparison figures are drawn from the original *AI OrgReview* (Sia, September 2024).

/ World Economic Forum, Future of Jobs Report, 2025.

/ OpenAI GPT-4 at launch (March 2023): ~$30 per million input tokens. Anthropic Claude Opus 3 (2024): $15 per million input tokens; Claude Opus 4.5–4.7 (2025–2026): $5 per million input tokens. Sub-frontier models (e.g., Claude Haiku 4.5, GPT-4o mini): $0.15–0.50 per million input tokens. Sources: Anthropic pricing documentation (platform.claude.com); OpenAI published pricing history.

/ McKinsey & Company, The State of AI, 2025.

/ Stanford Digital Economy Lab, *Generative AI at Work*, 2025 (+15% productivity on issues resolved per hour). MIT economics study (*economics.mit.edu*, 2025): +26% completed tasks for software developers.

/ Sources underlying sector benchmark table: Gartner Hype Cycle for AI in Finance 2025; Gartner AI Agents: Case Examples from Banking (Dec 2025); Gartner AI Agents: Case Examples for P&C and Life Insurance (Oct 2025); Gartner Software Engineering AI Briefing (Feb 2026); NVIDIA Industry Use Cases 2024; MIT productivity study 2025; Sia client engagement data and SME benchmark 2025–2026.

Optimists for change

Sia is a next-generation, global management consulting group—born digital, augmented by data, enhanced by creativity, and driven by responsibility. We partner with clients to resolve challenges and capitalize on opportunities. We believe that in today's world of change and disruption, optimism is a force multiplier.