Proceedings of the 7th Latin American Conference on Industrial Engineering and Operations Management
Panama City, Panama, August 4-6, 2026

Proceedings of the International Conference on Industrial Engineering and Operations Management

Publisher: IEOM Society International, USA DOI: 10.46254/SA07.20260128
Published: August 4, 2026

The AI Budget Paradox: A Tiered Access and Proficiency Framework for Sustainable Enterprise AI Adoption

Kirandeep Virdi
New York University
New York, NY, USA
kv2433@nyu.edu

Sergio Mastrogiovanni
New York University, New York, NY, USA
IAE Business School, Buenos Aires, Argentina
sm7758@nyu.edu, smastrogiovanni@alumni.iae.edu.ar

Abstract

Consumption-based pricing for enterprise artificial intelligence (AI) coding tools has broken the seat-based budgeting assumptions on which enterprise software procurement has long relied. Documented 2026 cases illustrate the mismatch: Uber exhausted its annual AI coding budget four months into the year, Microsoft began cancelling internal AI coding licences weeks later, and an Nvidia executive reported that AI compute now costs more than the engineers who use it. This paper reframes these events not as failures of AI, but as evidence of a structural governance gap: adoption succeeded while cost governance did not. Using a multiple-case synthesis of secondary evidence combined with a design-science construction of a management framework, the paper proposes a Tiered Access and Proficiency Framework that couples three disciplines already present in most enterprises—financial operations (FinOps), skills development, and security governance—into a single system in which demonstrated proficiency determines access, access determines cost, and cost is monitored by existing FinOps infrastructure within boundaries set by recognised governance standards. An illustrative financial model for a 2,000-engineer organisation estimates a 42% first-year and 51% steady-state cost reduction while retaining approximately 87% of productivity value, with the value-to-cost ratio improving as the workforce matures. The framework is positioned as an integration layer rather than new bureaucracy, mapping proficiency to the Skills Framework for the Information Age (SFIA) and governance to the NIST AI Risk Management Framework and ISO/IEC 42001. Limitations, applicability boundaries, and implementation risks are discussed.

Keywords
Enterprise AI governance, FinOps, AI proficiency, cost management, and operations management.

1. Introduction
For most of the history of enterprise software, tools were priced per seat. An organisation purchased a licence for each employee, cost was fixed, and the annual budget was predictable. Generative AI coding tools broke that model. They are priced by consumption—by the volume of tokens an engineer generates through interaction with the model—so cost scales with use. Because these tools are genuinely useful, engineers use them heavily, and spending grows in proportion to adoption rather than to headcount. The result is a category of enterprise cost that existing budgeting, governance, and organisational design were not built to absorb.

The consequences became visible in early 2026. At Uber, AI coding assistants were deployed to engineers in late 2025; within three months a large majority of developers were classified as agentic coding users and a substantial

© IEOM Society International

share of committed code originated from AI systems, with monthly spending ranging from roughly US$150 to US$250 per engineer on average and heavy users consuming US$500 to US$2,000 per month (Pragmatic Engineer 2026; Forbes 2026). The annual budget was exhausted by April 2026, four months into the year. Microsoft began cancelling internal AI coding licences weeks later (The Verge 2026), and an Nvidia executive acknowledged publicly that compute costs for an internal team now exceed the salaries of the engineers using them (Axios 2026). These are not isolated anecdotes; they are early signals of a structural mismatch that affects any organisation deploying consumption-priced AI tools at scale. Survey evidence indicates that 42% of organisations report only partial visibility into AI spending and nearly half have scaled back AI agent deployments because costs outweighed benefits (KPMG 2026), even as aggregate AI spending continues to grow sharply (Gartner 2026).

The problem is fundamentally one of operations management rather than technology. The budgeting model assumes fixed costs while the pricing model charges by consumption; as adoption accelerates, costs grow non-linearly while budgets remain linear. This paper argues that the appropriate response is neither to restrict AI use nor to accept uncontrolled spend, but to govern the relationship between who receives the most expensive access and who is demonstrably equipped to generate value from it.

1.1 Objectives

This paper pursues four objectives. First, to characterise the AI budget paradox as a structural cost-governance problem arising from the collision of consumption pricing with seat-based budgeting. Second, to synthesise the available productivity evidence and show why unconditioned access maximises cost for uncertain value. Third, to construct a management framework—the Tiered Access and Proficiency Framework—that links proficiency, access, and cost within recognised governance standards without introducing new bureaucracy. Fourth, to evaluate the framework through an illustrative financial model and to specify its applicability boundaries, implementation path, and risks. The central contribution is an integration layer that wires together cost monitoring, skills development, and security governance into a single, measurable operating model.

2. Literature Review

The literature relevant to this framework spans four streams: the productivity evidence on generative AI, the economics of consumption pricing and financial operations, skills and proficiency frameworks, and AI governance standards. Each stream is mature in isolation, yet they remain largely disconnected in practice—a gap this paper addresses directly.

The productivity stream provides both encouragement and caution. Controlled and field studies report meaningful task-level gains from AI assistance: a field experiment on customer-support agents found productivity improvements concentrated among less-experienced workers (Brynjolfsson et al. 2023), and a controlled study of software developers using an AI pair-programmer reported substantial reductions in task completion time (Peng et al. 2023). Task-exposure analyses suggest that a large share of occupations are affected by large language models to some degree (Eloundou et al. 2023). Practitioner datasets echo these gains—across a large developer population the average engineer is reported to save several hours per week (DX 2026). Yet the enterprise-level picture is more ambiguous. Perceived improvements frequently exceed measured improvements: a randomised controlled trial found that experienced developers were 19% slower when using AI assistance despite believing they were 20% faster (METR 2025), and executive surveys document perceived productivity gains that exceed measured gains (Baslandze et al. 2026). A survey of nearly 6,000 executives across four countries finds that while 69% of firms use AI, roughly nine in ten report no detectable firm-level impact on productivity or employment to date (Yotzov et al. 2026). This gap between task-level and firm-level value is the empirical justification for conditioning access on proficiency and for redesigning roles rather than simply distributing tools.

The second stream concerns consumption pricing and financial operations. The FinOps discipline emerged to govern variable cloud spending by attributing consumption to teams and workloads and by embedding cost accountability in engineering decisions (FinOps Foundation 2026). A specialised set of AI-cost tools now extends this practice to token consumption (CloudZero 2026). The limitation is temporal: FinOps instrumentation reports consumption after it occurs and does not, by itself, prevent overspend. This motivates a control that acts before consumption—namely, the pricing tier to which a user is assigned.

© IEOM Society International

The third stream is skills measurement. The Skills Framework for the Information Age (SFIA) is the most widely adopted global standard for classifying information-technology competencies and levels of responsibility (SFIA Foundation 2026). Emerging AI-specific proficiency instruments extend this tradition (DataCamp 2026), but existing proficiency platforms measure skill without triggering changes to tool access. The absence of a provisioning link between proficiency measurement and access governance is a recurring gap.

The fourth stream is AI governance. The NIST AI Risk Management Framework provides a voluntary structure for identifying and managing AI risk (NIST 2023), and ISO/IEC 42001 establishes requirements for an AI management system (ISO 2023). These instruments classify and manage risk at the system level but do not address user-level cost governance. Complementary workforce evidence indicates rising AI adoption alongside falling worker confidence and elevated retention anxiety (ManpowerGroup 2026), and analyses of task automatability emphasise that human involvement remains essential to capturing value (Cognizant 2024). In summary, each stream supplies a necessary component—productivity rationale, cost instrumentation, skills measurement, and risk governance—but none connects them into a single mechanism in which proficiency conditions access and access conditions cost. The framework proposed here is that connective layer.

3. Methods

The study adopts a qualitative, design-oriented research strategy appropriate to an emerging phenomenon for which large-sample primary data does not yet exist. Two complementary methods are combined. The first is a multiple-case synthesis of publicly documented enterprise experiences with consumption-priced AI coding tools, used to establish the existence, mechanism, and boundary conditions of the cost-governance problem. Cases were selected purposively for information richness and documentation quality rather than representativeness, following established guidance for theory-building from case research in operations and management (Eisenhardt 1989; Yin 2018).

The second method is design science, in which the primary output is an artefact—here, a management framework—constructed to solve a defined class of problem and evaluated against stated objectives (Hevner et al. 2004; Peffers et al. 2007). Framework construction proceeded by identifying the independent disciplines already present in enterprises (FinOps, skills frameworks, and governance standards), specifying the missing couplings between them, and designing an integration mechanism that requires no new infrastructure. Evaluation is performed through an illustrative financial model that translates the framework into a cost and productivity comparison for a representative organisation, together with a structured analysis of applicability, implementation, and risk. The model is deliberately parametric and illustrative; its role is to demonstrate directional magnitude and internal logic, not to forecast a specific firm's outcome.

4. Data Collection

Evidence was drawn from three source categories. Case evidence on enterprise cost overruns and adoption intensity was compiled from documented industry reporting during the first half of 2026 (Pragmatic Engineer 2026; Forbes 2026; The Verge 2026; Axios 2026). Productivity and workforce evidence was compiled from practitioner datasets and survey research (DX 2026; METR 2025; Yotzov et al. 2026; Baslandze et al. 2026; ManpowerGroup 2026; Cognizant 2024) and from peer-reviewed and working-paper studies of generative-AI productivity (Brynjolfsson et al. 2023; Peng et al. 2023; Eloundou et al. 2023). Framework and standards evidence was drawn from the reference documentation of the cited skills and governance instruments (SFIA Foundation 2026; NIST 2023; ISO 2023; FinOps Foundation 2026; CloudZero 2026).

Model parameters—tier distribution, per-tier pricing, governance costs, and productivity retention—were derived from the ranges reported in the case and pricing evidence and are stated explicitly so that any organisation can recalibrate them against its own usage data. Because several sources are contemporary industry publications rather than archival records, all quantitative claims are treated as indicative and are reported with their originating source. This is a recognised limitation of researching a fast-moving phenomenon and is addressed further in Section 6.

5. Results and Discussion

5.1 The Tiered Access and Proficiency Framework

The framework rests on a single principle: AI tool access should be matched to demonstrated proficiency and role need rather than distributed by default. It operationalises this principle through three interconnected mechanisms—tiered access, proficiency gating, and an integration layer—illustrated in Figure 1.

© IEOM Society International

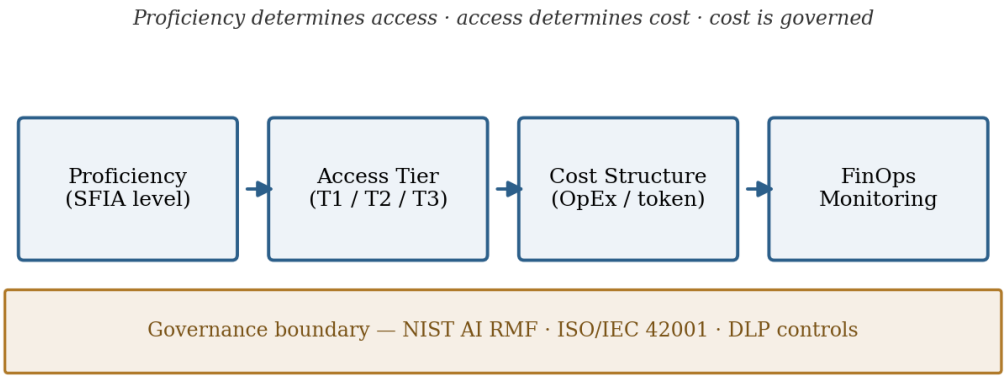

Figure 1. The Tiered Access and Proficiency Framework as an integration layer

The workforce is segmented into three access tiers, each with a distinct pricing strategy and governance posture. Tier 1 (must-have) comprises engineers in production-critical roles with measurable AI return on investment—backend, platform, and machine-learning engineers whose agentic workflows drive both the highest value and the highest token costs (approximately US$150–US$500 per month). They receive enterprise-grade access with full data-protection agreements, monitored by FinOps tooling. Tier 2 (good-to-have) is where the majority of the workforce begins. These users receive retail AI subscriptions (approximately US$10–US$60 per month) reimbursed through the learning-and-development (L&D) budget. The critical design insight is structural: retail plans carry built-in rate limits that function as automatic spending caps, so the pricing structure itself—rather than organisational discipline—constrains consumption. Tier 2 users are restricted to non-proprietary work, enforced technically by data-loss-prevention (DLP) controls that prevent proprietary code from reaching consumer AI platforms. Tier 3 (investment) is a small, selected group granted uncapped enterprise access for experimentation, with full outcome tracking and every token attributed to a project (approximately US$400–US$1,500+ per month).
A distinctive element is the repurposing of existing L&D budgets as the primary vehicle for adoption at scale. Rather than negotiating enterprise-wide licences upfront, organisations use L&D reimbursement to place AI tools in every employee's hands immediately, at retail pricing with built-in caps, while building the proficiency that determines who graduates to production-grade access. This aligns individual skill development with organisational cost governance in a way neither function achieves alone.

5.2 Proficiency Gating and the Integration Layer

Tier placement is neither permanent nor determined by job title alone. It is governed by a dual gate: the role determines the maximum tier for which an engineer is eligible, and a proficiency assessment determines the tier actually granted. The proficiency model maps to SFIA rather than to a bespoke scheme, progressing from basic AI awareness (default Tier 2 entry) through capable prompting to proficient multi-step workflow management with token-cost awareness (qualifying for Tier 1) and architect-level pipeline design (Tier 3). Assessments are on-demand and recalibrated periodically, because the competency frontier shifts rapidly. When an engineer passes the qualifying assessment and the role qualifies, provisioning into enterprise access is automatic, removing the procurement friction that would otherwise drive proficient engineers toward unsanctioned workarounds.

Critically, the framework replaces nothing. AI FinOps tools already track token spend but monitor after the fact; proficiency frameworks already measure skills but do not trigger provisioning; governance standards already classify risk but do not address user-level cost. The framework is the wiring between them: proficiency level determines access tier, which determines cost structure, which is monitored by FinOps, within governance boundaries set by risk classification. It is an integration layer, not a new bureaucracy

5.3 Numerical Results

Consider a mid-sized technology organisation with approximately 2,000 engineers. Under an unmanaged deployment—every engineer receiving enterprise access without tiering or gating—the budgeted annual cost is approximately US$8.4 million at a vendor-averaged US$350 per engineer per month. As the case evidence

© IEOM Society International

demonstrates, that figure is unreliable: when adoption accelerates without caps, realistic exposure runs three to four times higher, to roughly US$25–US$34 million. Under the tiered framework, with an illustrative distribution of 30% of engineers in Tier 1, 65% in Tier 2, and 5% in Tier 3—with Tier 1 consumption modelled at US$250 per engineer per month, within the documented US$150–US$500 range, reflecting FinOps-monitored rather than unmanaged usage—tool cost falls to approximately US$3.8 million; adding an assessment platform and incremental DLP controls brings the total to roughly US$4.84 million in Year 1 and US$4.09 million from Year 2 onward. Table 1 summarises the tiered cost structure.

Table 1. Illustrative tiered cost structure for a 2,000-engineer organization

Tier (share, count)	Plan	Monthly (US$)	Annual (US$)
Tier 1 (30%, 600)	Enterprise (seat + API)	250	1,800,000
Tier 2 (65%, 1,300)	Retail (rate-limited)	35	546,000
Tier 3 (5%, 100)	Enterprise, uncapped	1,200	1,440,000
Governance (assessment + DLP), Year 1	—	—	1,050,000
Total, Year 1	—	—	4,836,000
Total, Year 2 onward	—	—	4,086,000

Relative to the budgeted unmanaged scenario, the tiered model yields an estimated 42% saving in Year 1 and 51% from Year 2; relative to the realistic overrun, the saving approaches 80–85%. Figure 2 contrasts the four cost positions.

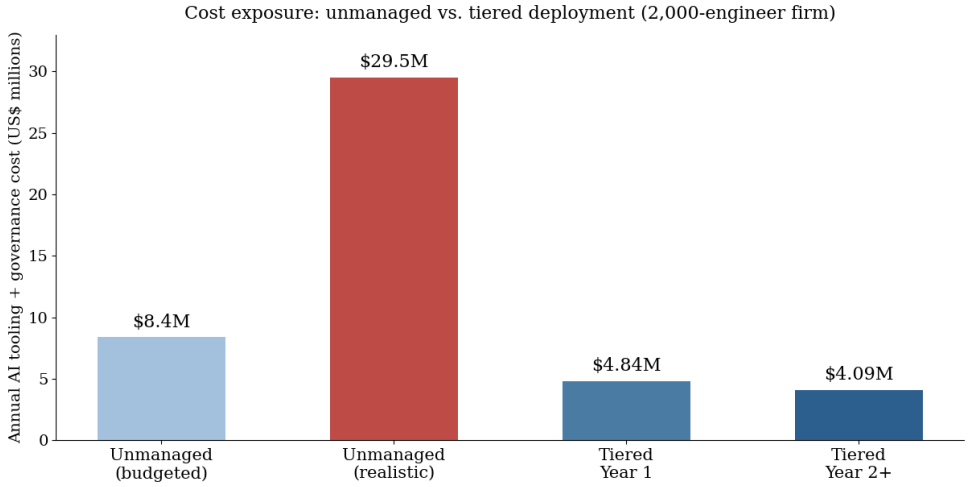

Figure 2. Estimated annual cost exposure under unmanaged and tiered deployment

Productivity requires honesty. The tiered approach retains an estimated 87% of the productivity value of unmanaged deployment; the 13% reduction arises primarily from Tier 2 engineers working on non-proprietary tasks at lower intensity during skill-building. This gap closes over time. As roughly 15% of Tier 2 engineers graduate to Tier 1 each year, productive capacity grows while governance remains intact; by Year 3, with a distribution shifting toward approximately 55% Tier 1, 35% Tier 2, and 10% Tier 3, estimated net productivity can exceed the unmanaged scenario at materially lower cost. The framework therefore does not save money by restricting AI use; it saves money by ensuring that the most expensive access reaches the people demonstrably equipped to generate value from it, tightening the spend-to-value ratio.

5.4 Organisational Implications and Applicability

Two organisational changes accompany the cost mechanism. First, token costs should be classified as business infrastructure—the same operating-expenditure category as cloud compute—rather than loaded into cost-of-hire, preserving recruiting metrics while creating a governable budget line. Cost-of-hire changes not because tokens are added to it, but because roles themselves transform toward AI-orchestration profiles that already command distinct

© IEOM Society International

market rates. Second, the programme must be framed as career development, not access restriction: each proficiency level is a credential mapped to SFIA and linked to progression, and L&D reimbursement is tangible evidence of investment. Given documented declines in workforce confidence (ManpowerGroup 2026), this framing is not cosmetic but material to adoption. Role evaluation must run in parallel, because optimising access without redesigning work risks accelerating misaligned tasks—precisely the firm-level productivity paradox the evidence describes (Yotzov et al. 2026; Cognizant 2024). Applicability depends on three variables: whether AI tools are consumption-priced rather than embedded, whether the engineering population is large enough for governance overhead to be proportionate, and whether the operating environment has mature enterprise IT infrastructure. The core fit is technology and software firms with roughly 500–10,000 engineers in advanced digital economies. Financial and professional services fit with adaptation, often requiring higher Tier 1 allocation for regulatory reasons. Notably, in high-growth digital economies—including much of Latin America—where tool prices are global but engineer salaries are markedly lower, the token-to-salary ratio is more acute, making cost governance proportionally more valuable rather than less. Limited fit applies to sectors where AI is embedded in operational systems rather than consumed as coding tools, and to organisations below roughly 200 engineers, where a simplified two-tier model with an acceptable-use policy captures most of the value at a fraction of the governance cost.

5.5 Sensitivity Analysis and Validation

Because the financial model is illustrative, its conclusions should hold under reasonable perturbation of its parameters. Varying the steady-state tier distribution from the baseline toward a more expensive profile (40% Tier 1, 55% Tier 2, 5% Tier 3) raises the steady-state total to approximately US$4.60 million and reduces the saving from 51% to approximately 45%, while a more conservative profile (20% Tier 1, 75% Tier 2, 5% Tier 3) lowers the total to approximately US$3.57 million and raises the saving to approximately 58%. Raising the modelled Tier 1 price by 40%, from US$250 to US$350 per engineer per month, still yields an estimated 43% steady-state saving, and raising the Tier 3 allowance from US$1,200 to US$1,500 per month yields approximately 47%. Across all perturbations, the saving relative to the budgeted unmanaged scenario remains between roughly 43% and 58%, and the saving relative to the realistic overrun scenario remains above 80%. The direction of the result—material cost reduction with retained productivity—is therefore robust to the parameter choices an adopting organisation is most likely to change. Validation of the framework itself is necessarily staged. Internal validation is provided by the consistency of the model parameters with documented pricing and adoption ranges and by the robustness analysis above. External validation—longitudinal measurement of cost, proficiency progression, and productivity in organisations that adopt the framework—is defined as future work and is a precondition for any stronger empirical claim.

6. Conclusion

The AI budget overruns documented at leading technology firms are not evidence of AI failing; they are evidence of organisations succeeding at adoption without succeeding at governance. The tools delivered what they promised, and the cost of that delivery exceeded every model because the models were built for fixed-price software, not consumption-priced intelligence. This paper set out to characterise that mismatch, to reconcile the productivity evidence with it, to construct a framework that resolves it, and to evaluate the framework's economics and boundaries—the four objectives stated in Section 1.1, each of which the analysis has addressed.

The unique contribution is an integration layer that links proficiency to access, access to cost, and cost to FinOps monitoring within recognised governance standards, using infrastructure organisations already possess. Its value compounds: as engineers develop proficiency and graduate through tiers, as roles are redesigned around AI capability, and as the organisation accumulates a dataset of who can do what with AI at what cost and to what effect, the framework becomes not merely a cost-control tool but a strategic asset. The principal limitations are the illustrative nature of the financial model and the contemporary, fast-moving character of several evidentiary sources; both point to the same future work—empirical calibration and longitudinal testing of the framework within organisations that adopt it. The organisations that navigate this era well will not be those that spend the most on AI, but those that know which spending produces value, who is equipped to generate it, and whether the work being done still makes sense.

References

Axios, AI can cost more than human workers now, Available: https://www.axios.com/2026/04/26/ai-cost-human-workers, Accessed on July 14, 2026.

© IEOM Society International

Baslandze, S., Edwards, Z., Graham, J., McClure, T., Meyer, B. H., Sparks, M., Waddell, S. R. and Weitz, D., Artificial intelligence, productivity, and the workforce: evidence from corporate executives, National Bureau of Economic Research, Working Paper No. 34984, 2026, https://doi.org/10.3386/w34984.

Brynjolfsson, E., Li, D. and Raymond, L. R., Generative AI at work, National Bureau of Economic Research, Working Paper No. 31161, 2023, https://doi.org/10.3386/w31161.

CloudZero, How much does AI cost? The complete guide for 2026, Available: https://www.cloudzero.com/blog/how-much-does-ai-cost/, Accessed on July 14, 2026.

Cognizant, New work, new world, Cognizant and Oxford Economics, 2024. Available: https://www.cognizant.com/en_us/insights/documents/new-work-new-world-with-generative-ai-wf2064768.pdf, Accessed on July 14, 2026.

DataCamp, The 2026 state of data and AI literacy report, DataCamp and YouGov, 2026. Available: https://www.datacamp.com/blog/the-state-of-data-and-ai-literacy-in-2026-definitions-statistics-and-the-ai-skills-gap, Accessed on July 14, 2026.

DX, Measuring the impact of AI coding tools on developer productivity, DX Research, 2026. Available: https://getdx.com/, Accessed on July 14, 2026.

Eisenhardt, K. M., Building theories from case study research, *Academy of Management Review*, vol. 14, no. 4, pp. 532-550, 1989.

Eloundou, T., Manning, S., Mishkin, P. and Rock, D., GPTs are GPTs: an early look at the labor market impact potential of large language models, arXiv:2303.10130, 2023.

FinOps Foundation, State of FinOps report, FinOps Foundation, 2026. Available: https://www.finops.org/, Accessed on July 14, 2026.

Forbes, Uber burns its 2026 AI budget in four months on Claude Code, Forbes, May 17, 2026. Available: https://www.forbes.com/sites/janakirammsv/2026/05/17/uber-burns-its-2026-ai-budget-in-four-months-on-claude-code/, Accessed on July 14, 2026.

Gartner, Gartner forecasts worldwide AI spending to grow 47% in 2026, Gartner Newsroom, May 19, 2026. Available: https://www.gartner.com/en/newsroom/press-releases/2026-05-19-gartner-forecasts-worldwide-ai-spending-to-grow-47-percent-in-2026, Accessed on July 14, 2026.

Hevner, A. R., March, S. T., Park, J. and Ram, S., Design science in information systems research, *MIS Quarterly*, vol. 28, no. 1, pp. 75-105, 2004.

International Organization for Standardization, ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system, ISO, Geneva, 2023.

KPMG, Global AI pulse: Q2 2026, KPMG International, July 2026.

ManpowerGroup, 2026 Global Talent Barometer, ManpowerGroup, 2026.

METR, Measuring the impact of early-2025 AI on experienced open-source developer productivity, METR Research, July 2025. Available: https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/, Accessed on July 14, 2026.

National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, U.S. Department of Commerce, 2023, https://doi.org/10.6028/NIST.AI.100-1.

Peffers, K., Tuunanen, T., Rothenberger, M. A. and Chatterjee, S., A design science research methodology for information systems research, *Journal of Management Information Systems*, vol. 24, no. 3, pp. 45-77, 2007.

Peng, S., Kalliamvakou, E., Cihon, P. and Demirer, M., The impact of AI on developer productivity: evidence from GitHub Copilot, arXiv:2302.06590, 2023.

Pragmatic Engineer, How Uber uses AI for development: inside look, The Pragmatic Engineer, March 2026. Available: https://newsletter.pragmaticengineer.com/p/how-uber-uses-ai-for-development, Accessed on July 14, 2026.

SFIA Foundation, Skills Framework for the Information Age (SFIA), SFIA Foundation, 2026. Available: https://sfia-online.org/, Accessed on July 14, 2026.

The Verge, Microsoft cancels internal Claude Code licences (Notepad newsletter by T. Warren), The Verge, May 14, 2026.

Yin, R. K., *Case Study Research and Applications: Design and Methods*, 6th Edition, Sage Publications, Thousand Oaks, CA, 2018.

Yotzov, I., Barrero, J. M., Bloom, N., Bunn, P., Davis, S. J., Foster, K. M., Jalca, A., Meyer, B. H., Mizen, P., Navarrete, M. A., Smietanka, P., Thwaites, G. and Wang, B. Z., Firm data on AI, National Bureau of Economic Research, Working Paper No. 34836, 2026, https://doi.org/10.3386/w34836.

© IEOM Society International

Biographies

Kirandeep Virdi is a human resources strategy, technology, and transformation leader affiliated with New York University, where her work centers on the workforce implications of enterprise artificial intelligence adoption. Across a career in people-and-culture leadership at high-growth technology organizations, she has directed talent management, HR technology, and organizational transformation programs. She has served as Head of People and Culture at Asper.ai, an enterprise artificial-intelligence and decision-intelligence company; as Associate Director for Talent and Diversity & Inclusion at Sprinklr; and in senior human-resources roles within Airtel's global-business and data-centre operations. Throughout these positions she has focused on building talent architectures for technology-intensive firms—designing capability frameworks, aligning human-resources systems with business strategy, and redesigning roles around emerging digital and AI competencies. Virdi holds a postgraduate qualification in Human Resource Management and Labour Relations from the Tata Institute of Social Sciences (TISS), among the leading programs of its kind in India. Her current research examines the intersection of human capital and enterprise AI governance: specifically, how organizations can match AI-tool access to demonstrated proficiency, embed workforce development within cost-governance mechanisms, and convert fragmented adoption into sustainable capability. This emphasis on proficiency-based provisioning and the redesign of roles around AI capability directly informs the Tiered Access and Proficiency Framework advanced in the present paper.

Sergio Mastrogiovanni is a Postdoctoral Researcher at IAE Business School, Universidad Austral, Buenos Aires, Argentina, and an Adjunct Professor at New York University, with teaching appointments at IAE Business School and Universidad Austral. A global executive in artificial intelligence and digital strategy, he brings more than three decades of experience leading the convergence of data, business, and public policy, having designed and implemented transformation programs for Fortune 500 companies—including multinationals such as Pfizer—as well as startups and government entities. He is the co-founder and principal consultant of Socradata, a boutique artificial-intelligence and advanced-analytics firm based in Buenos Aires that drives transformation across enterprises and smart cities. His research examines enterprise AI governance, intelligent automation, data science, human-centred smart cities, digital identity, and data-driven urban management, integrating academic inquiry with applied consulting practice. His applied work includes blockchain-based digital-identity deployments with the Government of the City of Buenos Aires (QuarkID) and computer-vision and AI applications in airport logistics and urban mobility, advancing transparency and citizen participation in Buenos Aires and New York City. He holds a Ph.D. in Smart Cities and Artificial Intelligence from IAE Business School, an MBA from the NYU Stern School of Business, and postgraduate credentials in AI and data from MIT, Columbia University, and Microsoft/AWS. His work has been accepted for presentation at the Academy of Management Annual Meeting (2026).

© IEOM Society International