Agentic AI in Enterprise Business Processes: A Systematic Review and Practitioner Survey on Adoption Readiness

Sergey Emelyanov
Graduate School of Business
HSE University
Moscow, Russia
sbemelyanov@edu.hse.ru

https://doi.org/10.63962/NXPA6137

Abstract—Agentic AI systems—large language model-based agents capable of autonomous reasoning, planning, and tool use—are entering enterprise workflows at an accelerating pace. This study combines a systematic review of 47 academic papers and 12 industry reports (2023–2025) with an original cross-sectional survey of 136 IT professionals across 14 industries. The review identifies six dominant architectural patterns and a persistent gap between vendor-reported and independently verified outcomes. The survey reveals that 73% of respondents consider organizational readiness—not technical capability—the primary barrier to agentic AI adoption. Based on both evidence streams, we propose and empirically ground the Enterprise Agentic AI Maturity Model (EAAMM), a five-level framework integrating technical autonomy with organizational readiness.

Keywords—agentic AI, autonomous agents, large language models, enterprise automation, systematic review, practitioner survey, maturity model

I. Introduction

Large language models have evolved from conversational tools into agents that pursue goals through multi-step reasoning, tool invocation, and coordinated problem-solving [1], [2]. What distinguishes an agentic system is its capacity for iterative, goal-directed behavior: the agent perceives its environment, formulates a plan, selects tools, executes actions, and revises its approach based on feedback [3]. The ReAct framework, which interleaves reasoning traces with action steps, has become the architectural template for most enterprise deployments [1].

Adoption figures reflect this shift. McKinsey's 2025 global survey (n=1,993) found 88% of organizations using AI regularly, 62% experimenting with AI agents, and 23% having scaled agentic AI in at least one function [4]. Industry estimates project compound annual growth rates above 40% through 2030, with Gartner forecasting that 33% of enterprise software will embed agentic AI by 2028 [5], [6]. Yet Gartner also projects that over 40% of agentic AI projects will be canceled by late 2027 owing to cost overruns and insufficient risk controls [6]. Reported adoption varies by sector: financial services and technology lead at 35–40%, while healthcare and manufacturing lag at 15–20%, though these figures rely primarily on vendor-commissioned surveys [4], [14]. IDC projects significant growth in AI agent deployment through 2029 [13].

This paradox—strong investment momentum coupled with high failure rates—motivates our study. Existing literature offers either technical architectural surveys [1]–[3] or vendor-commissioned adoption reports [4]–[6], but lacks empirical practitioner perspectives on readiness barriers. This paper addresses the gap through three contributions: (1) a systematic review mapping the agentic AI landscape, (2) an original practitioner survey (n=136) identifying adoption barriers and self-assessed maturity levels, and (3) the Enterprise Agentic AI Maturity Model (EAAMM), empirically grounded through both evidence streams.

II. Methodology

A. Systematic Review Protocol

The review followed PRISMA-ScR guidelines [9]. Searches spanned IEEE Xplore, ACM Digital Library, arXiv, Scopus, and Google Scholar using the query: (agentic AI OR AI agents OR autonomous agents OR LLM agents) AND (enterprise OR business process OR workflow OR automation), limited to January 2023 through December 2025. The initial search returned 312 candidates; title and abstract screening narrowed this to 127, and full-text review yielded 47 academic papers and 12 industry reports from Gartner, McKinsey, Deloitte, Forrester, IDC, and Capgemini. Data extraction separated peer-reviewed academic findings from vendor-commissioned and consulting reports, flagging vendor-reported figures distinctly from independently validated results to address persistent reporting bias in enterprise AI literature.

B. Practitioner Survey Design

To complement the literature review with primary empirical data, a cross-sectional online survey was conducted in January 2026 using a structured questionnaire distributed via convenience sampling through LinkedIn professional groups, Telegram communities, and direct outreach to

XXX-X-XXXX-XXXX-X/XX/$XX.00 ©20XX IEEE

enterprise technology teams. The instrument comprised 22 items across four sections: (1) organizational demographics, (2) current AI adoption status and self-assessed maturity level, (3) perceived barriers to agentic AI deployment ranked on a five-point Likert scale, and (4) governance readiness indicators. Of 198 responses received, 136 met inclusion criteria (respondents with direct involvement in AI strategy or implementation at organizations with 50+ employees). The sample spans 14 industry sectors and 9 countries, with technology (27%), financial services (19%), manufacturing (14%), and retail (11%) as the largest cohorts. Respondent roles include product managers (26%), engineering leads (22%), CTOs and VPs of technology (17%), AI/ML engineers (16%), and business analysts (12%). Company size distribution: under 500 employees (29%), 500–5,000 (34%), and over 5,000 (37%). The inclusion rate was 68.7% (136 of 198). The sample is geographically concentrated, with Russian respondents comprising 35.3%; however, cross-tabulation shows that Russian response patterns do not differ significantly from the overall sample on key variables (governance status, barrier rankings). Internal consistency was verified with Cronbach's alpha of 0.81 for the barrier perception scale [26].

III. Preliminary Results

A. Architectural Patterns

Since 2023, the architectural landscape has coalesced around well-defined patterns. ReAct [1] remains the standard for single-agent systems, and Reflexion extends it through verbal self-critique [25]. For complex processes, multi-agent architectures dominate [8]: AutoGen [7] uses conversation-based protocols, MetaGPT [10] assigns software-engineering roles, and CrewAI organizes hierarchical teams. Six orchestration patterns recur: sequential chaining, parallel fan-out, routing-based delegation, hierarchical manager–worker coordination, peer-to-peer negotiation, and evaluator–optimizer feedback loops [11], [12]. Our analysis indicates that roughly 80% of current enterprise installations rely on workflow-based orchestration rather than full autonomy [11], [14]. A critical architectural consideration is memory and context management: long-term memory mechanisms including retrieval-augmented generation (RAG) and vector databases enable agents to maintain state across sessions [20], [21], but also introduce failure modes such as stale context, hallucinated recollections, and privacy risks from accumulated data [3], [8].

B. Enterprise Adoption Evidence

Data from major industry surveys reveal broad experimentation paired with limited production rollout [4]–[6], [13], [14]. Between 60% and 88% of enterprises use AI, yet only 2–14% have deployed agentic AI at production scale. Capgemini's 2025 survey finds 2% full-scale deployment, 12% partial, 23% active pilots, and 61% still exploring [14]; a separate 2024 study found 82% plan agentic AI integration within one to three years [22]. Deloitte's 2026 edition notes that 74% intend to deploy agents within two years, but only 21% have the governance infrastructure to support it [5]. Vendor case studies present impressive figures: Salesforce reported 84% autonomous resolution across 380,000 customer support conversations [15]; Klarna's AI assistant handled 2.3 million chats in its first month [16]. However, these narratives face contradictory independent evidence. An observational study of approximately 800 developers found no throughput gains and 41% more bugs with AI coding assistants [23], while a randomized trial with 16 developers found AI tools made them 19% slower [24]. Across independent surveys, only 15% of AI decision-makers report EBITDA improvement [18], and merely 6% qualify as AI high performers [4]. The customer service domain offers a particularly instructive pattern: initial deployments achieve high automation rates on straightforward inquiries, but as coverage expands to ambiguous cases, resolution quality declines and human escalation rates rise—a trajectory visible in Klarna's pivot to a hybrid model in 2025 and consistent with Gartner's prediction of widespread project cancellations [6], [16].

C. Practitioner Survey Findings

The survey produced five key findings that extend and contextualize the literature evidence.

First, self-assessed maturity levels reveal an early-stage landscape. When asked to classify their organization against the EAAMM framework, 41% placed themselves at L1 (Assistive), 32% at L2 (Collaborative), 19% at L3 (Consultative), 6% at L4 (Supervisory), and 2% at L5 (Autonomous). This distribution aligns with the 2–14% production deployment range found in the literature review, but provides finer granularity: most organizations are not merely "exploring" but are stuck at specific maturity transitions.

Second, organizational readiness dominates as the perceived primary barrier. When asked to identify the single greatest obstacle to advancing agentic AI maturity, 73% of respondents selected organizational factors—governance gaps (31%), data quality and infrastructure (22%), workforce skills (12%), and risk management concerns (8%)—versus 27% citing technical limitations such as model accuracy (15%), integration complexity (8%), or cost (4%). This finding is consistent with the hypothesis derived from the literature review that organizational readiness, not technical capability, constitutes the primary perceived constraint.

Third, a pronounced governance deficit was identified. Only 18% of respondents reported having formal AI governance frameworks in place, 35% had partial frameworks, and 48% had none. Cross-tabulation reveals a strong association between governance maturity and self-assessed EAAMM level: 89% of L3+ organizations have at least partial governance, compared to only 38% of L1 organizations ($\chi^2 = 27.4$, $p < 0.001$, Cramér's $V = 0.45$). This pattern is consistent with the EAAMM's gating mechanism, which posits that governance infrastructure is needed before advancing beyond L2.

Fourth, expected adoption timelines reveal cautious optimism. Among respondents not yet at L3, 14% expect to reach L3 within one year, 39% within one to two years, 31% within two to three years, and 16% consider it unlikely within three years. Notably, respondents from organizations with

established governance frameworks projected shorter timelines (mean 1.6 years) than those without (mean 2.8 years, $t(98)=3.42$, $p < 0.01$, Cohen's $d=0.69$), suggesting an association between governance readiness and shorter projected timelines. As this is cross-sectional data, the direction of causality cannot be established.

Fifth, formal ROI measurement remains rare. Only 21% of respondents have established frameworks for tracking agentic AI return on investment, while 44% rely on informal or anecdotal assessment and 35% perform no ROI tracking at all. This measurement deficit partially explains the disconnect between vendor-reported impact figures and independently observed outcomes: without systematic measurement, organizations lack the data to distinguish genuine productivity gains from displaced costs or reporting artifacts [17], [18].

TABLE I. PRACTITIONER SURVEY: KEY FINDINGS (N=136)

Metric	Result
Self-assessed at L1–L2	73%
Organizational barrier as primary	73%
No formal AI governance	48%
Governance → faster timeline	1.6 vs 2.8 years
Formal ROI tracking in place	21%

IV. PROPOSED EAAMM FRAMEWORK

TABLE II. ENTERPRISE AGENTIC AI MATURITY MODEL (EAAMM)

Level	Human Role	Key Characteristics
L1 Assistive	Operator	Explicit prompts; basic governance
L2 Collaborative	Collaborator	Suggests actions, human approves; data quality protocols
L3 Consultative	Consultant	Autonomous on routine tasks; formal governance required
L4 Supervisory	Approver	Independent operation; mature risk management
L5 Autonomous	Observer	Full autonomy within boundaries; audit trails

The EAAMM integrates technical autonomy levels from Feng et al. [19]—Operator, Collaborator, Consultant, Approver, Observer—with four organizational readiness dimensions: (1) governance infrastructure, (2) data architecture, (3) workforce adaptation, and (4) risk management. Twelve existing frameworks were reviewed [3], [6], [11], [12], [14], [19], [20]; none adequately bridges both technical autonomy and organizational readiness. The EAAMM fills this gap with a technology-agnostic, cross-industry model designed as a diagnostic tool rather than a linear progression: an organization may operate at L3 for customer service while remaining at L1 for financial decision support. Each maturity level prescribes specific prerequisites before advancement. For instance, progressing from L2 (Collaborative) to L3 (Consultative) requires established data quality protocols, formal escalation procedures, and demonstrated reliability metrics on routine tasks. This gating mechanism prevents organizations from deploying autonomous agents without adequate governance—a failure mode that Gartner identifies as a primary driver of project cancellations [6]. The model's four readiness dimensions can be calibrated to industry-specific regulatory requirements: in financial services the risk management dimension may carry elevated weight, whereas for technology companies data architecture and workforce adaptation become priorities.

The survey data provide empirical support for the EAAMM's core premise. The strong association between governance maturity and self-assessed EAAMM level ($\chi^2 = 27.4$, $p < 0.001$, Cramér's $V = 0.45$) suggests that the four readiness dimensions correspond to observable organizational differences. As an empirically informed diagnostic framework rather than a statistically validated predictive model, the EAAMM relies on self-reported maturity assessments that may reflect organizational perception rather than independently verified capability. Future research should validate the model longitudinally through repeated measurement and independent capability assessment.

V. DISCUSSION OF FINDINGS

Three cross-cutting insights emerge from the combined evidence. First, agentic AI suffers from a pronounced evidence gap: academic work concentrates on architectural innovation, while business impact evidence rests on vendor-funded metrics that are frequently methodologically opaque. The contradictions between vendor-reported outcomes and independent evaluations [23], [24] illustrate why standardized third-party assessment methodologies are urgently needed.

Second, the survey data confirm and extend the literature finding that organizational readiness is the binding constraint. The 73% figure for organizational barriers as primary closely mirrors Deloitte's finding that 74% plan agent deployment but only 21% have governance infrastructure [5]. This convergence across independent evidence sources—industry reports, academic reviews, and our practitioner survey—strengthens the finding's reliability and supports EAAMM's emphasis on organizational prerequisites.

Third, an emerging pattern across both literature and survey data is the convergence toward human-in-the-loop architectures as the pragmatic default for enterprise deployment. Even the most successful vendor implementations maintain human oversight for edge cases, and 73% of surveyed organizations self-assess at L1–L2. This suggests the industry is implicitly settling at L2–L3 rather than pursuing full L5 autonomy—a finding with direct implications for technology investment prioritization and organizational change management strategies.

VI. CONCLUSION

This study combines a systematic review of 47 academic papers and 12 industry reports with an original practitioner survey (n=136) to provide a multi-method assessment of agentic AI in enterprise settings. Three findings emerge. First, while technical foundations are increasingly solid with well-established architectural patterns, only 2–14% of organizations have reached production scale, and 73% of surveyed practitioners remain at L1–L2 maturity. Second, organizational readiness—not technical capability—is the binding constraint: 73% of respondents identified organizational barriers as primary, and the governance deficit

(48% with no formal AI governance) represents the most acute bottleneck. Third, governance readiness is associated with shorter projected adoption timelines, with governed organizations reporting estimates 43% shorter than ungoverned peers (Cohen's d=0.69).

The EAAMM provides a structured assessment tool integrating technical autonomy with organizational readiness across five deployment stages, empirically grounded through both literature synthesis and survey validation. For practitioners, the study suggests three recommendations: begin with workflow-based orchestration (L2–L3) rather than pursuing full autonomy; invest in governance infrastructure before scaling deployments; and establish independent measurement frameworks beyond vendor-supplied metrics. Limitations include the cross-sectional design, which supports associations but not causal claims; convenience sampling through professional networks; and geographic concentration (35.3% Russian respondents), although subgroup analysis indicates that Russian response patterns align with the overall sample on governance and barrier variables. Future research should prioritize longitudinal EAAMM validation through repeated measurement and independent capability audits, integration of regulatory compliance dimensions such as the EU AI Act, and replication across diverse cultural contexts.

Acknowledgment

The author thanks the anonymous reviewers for their constructive feedback and the survey participants for their time and insights.

References

[1] S. Yao et al., "ReAct: Synergizing reasoning and acting in language models," in Proc. ICLR, 2023.

[2] J. Schneider, "Generative to agentic AI: Survey, conceptualization, and challenges," arXiv:2504.18875, 2025.

[3] H. Abou Ali et al., "A comprehensive survey of agentic AI," Artif. Intell. Rev., vol. 59, art. 36, 2026, doi:10.1007/s10462-025-11148-z.

[4] McKinsey & Co., "The state of AI in 2025: Agents, innovation, and transformation," QuantumBlack Global Survey, Nov. 2025. [Online]. Available: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

[5] Deloitte AI Institute, "State of AI in the enterprise: The untapped edge," 2026 edition, Jan. 2026. [Online]. Available: https://www2.deloitte.com/content/dam/Deloitte/us/Documents/consulting/us-state-of-ai-in-the-enterprise.pdf

[6] Gartner, Inc., "Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027," Press Release, Jun. 2025. [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2025-06-agentic-ai-projects

[7] Q. Wu et al., "AutoGen: Enabling next-gen LLM applications via multi-agent conversation," in Proc. COLM, 2024.

[8] T. Guo et al., "Large language model based multi-agents: A survey of progress and challenges," in Proc. IJCAI, 2024.

[9] A. C. Tricco et al., "PRISMA extension for scoping reviews (PRISMA-ScR): Checklist and explanation," Ann. Intern. Med., vol. 169, no. 7, pp. 467–473, 2018.

[10] S. Hong et al., "MetaGPT: Meta programming for a multi-agent collaborative framework," in Proc. ICLR, 2024.

[11] Anthropic, "Building effective agents," Dec. 2024. [Online]. Available: https://www.anthropic.com/research/building-effective-agents

[12] Microsoft, "AI agent design patterns," Azure Architecture Center, 2025.

[13] IDC, "Agent economics: Adoption and delivery model," IDC Research (US54034725), Dec. 2025.

[14] Capgemini Research Inst., "Rise of agentic AI: How trust is the key to human-AI collaboration," Jul. 2025.

[15] Salesforce, Inc., "Fourth quarter and fiscal year 2025 results," Salesforce Investor Relations, Feb. 2025.

[16] Klarna, "Klarna AI assistant handles two-thirds of customer service chats in its first month," Press Release, Feb. 2024.

[17] Forrester Consulting, "The total economic impact of Microsoft 365 Copilot," commissioned by Microsoft, Mar. 2025.

[18] Forrester Research, "Predictions 2026: AI moves from hype to hard hat work," Oct. 2025.

[19] K. J. K. Feng, D. W. McDonald, and A. X. Zhang, "Levels of autonomy for AI agents," arXiv:2506.12469, Jun. 2025.

[20] T. R. Sumers, S. Yao, K. Narasimhan, and T. L. Griffiths, "Cognitive architectures for language agents," Trans. Mach. Learn. Res. (TMLR), 2024.

[21] L. Wang et al., "A survey on large language model based autonomous agents," Front. Comput. Sci., vol. 18, no. 6, art. 186345, 2024.

[22] Capgemini Research Inst., "Harnessing the value of generative AI," 2nd ed., Jul. 2024.

[23] Uplevel, "Can generative AI improve developer productivity?," Uplevel Data Labs Research Report, Sep. 2024.

[24] J. Becker et al., "Measuring the impact of early-2025 AI on experienced open-source developer productivity," arXiv:2507.09089, Jul. 2025.

[25] N. Shinn et al., "Reflexion: Language agents with verbal reinforcement learning," in Proc. NeurIPS, 2023.

[26] L. J. Cronbach, "Coefficient alpha and the internal structure of tests," Psychometrika, vol. 16, no. 3, pp. 297–334, 1951.

Supplementary Materials

Agentic AI in Enterprise Business Processes: A Systematic Review and Practitioner Survey on Adoption Readiness

Sergey Emelyanov, HSE University

APPENDIX A. PRACTITIONER SURVEY INSTRUMENT

The following questionnaire was administered online in January 2026 to IT professionals involved in AI strategy or implementation. The instrument comprises 22 items across four sections. Inclusion criteria: direct involvement in AI strategy or implementation at organizations with 50+ employees. Of 198 responses received, 136 met inclusion criteria. Internal consistency: Cronbach's α = 0.81 for the barrier perception scale (Section 3).

Section 1: Organizational Demographics (Q1–Q6)

Q1. In which industry does your organization primarily operate? *[Single choice, 14 options]*

Q2. How many employees does your organization have? *[Single choice: <500 / 500–5,000 / >5,000]*

Q3. In which country is your organization headquartered? *[Open text]*

Q4. What is your current role? *[Single choice: Product Manager / Engineering Lead / CTO or VP of Technology / AI-ML Engineer / Business Analyst / Data Scientist / Other]*

Q5. How many years of experience do you have in AI/technology roles? *[Numeric]*

Q6. Are you directly involved in AI strategy or implementation decisions at your organization? *[Yes / No (filter question)]*

Section 2: AI Adoption Status and Maturity (Q7–Q11)

Q7. Does your organization currently use AI (including traditional ML, generative AI, or agentic AI) in any business process? *[Yes, regularly / Yes, experimenting / No, but planning / No]*

Q8. What is the current status of agentic AI deployment at your organization? *[Single choice: Full-scale production / Partial production / Active pilot(s) / Exploring or evaluating / Not considering]*

Q9. Based on the following EAAMM framework descriptions, at which level would you classify your organization's most advanced agentic AI deployment? *[Single choice: L1 Assistive / L2 Collaborative / L3 Consultative / L4 Supervisory / L5 Autonomous]*

L1 Assistive: AI responds to explicit prompts; human operates and controls all actions. L2 Collaborative: AI suggests actions and drafts; human reviews and approves. L3 Consultative: AI handles routine tasks autonomously; human consulted on exceptions. L4 Supervisory: AI operates independently on complex tasks; human approves critical decisions. L5 Autonomous: AI operates with full autonomy within defined boundaries; human monitors.

Q10. What are the primary use cases for agentic AI at your organization? (Select all that apply) *[Multi-choice: Customer service / Software development / Data analysis / Document processing / Sales and marketing / Supply chain / HR and recruitment / Financial operations / Other]*

Q11. How many distinct agentic AI projects or initiatives does your organization currently have? *[Numeric: 0 / 1–2 / 3–5 / 6–10 / >10]*

Section 3: Perceived Barriers to Agentic AI Deployment (Q12–Q19)

Before rating each barrier individually, please indicate: Which of the following do you consider the single greatest barrier to advancing agentic AI maturity at your organization?

[Single choice: Governance gaps / Data quality and infrastructure / Workforce skills and training / Risk management concerns / Model accuracy and reliability / Integration complexity / Cost constraints / Regulatory compliance]

For each of the following factors, please rate the extent to which it represents a barrier to advancing agentic AI maturity at your organization.

Response scale: 1 = Not a barrier at all; 2 = Minor barrier; 3 = Moderate barrier; 4 = Significant barrier; 5 = Critical barrier

Q12. Governance gaps (lack of formal policies, oversight structures, accountability frameworks for AI agents) *[Likert 1–5]*

Q13. Data quality and infrastructure (insufficient data pipelines, poor data quality, inadequate integration architecture) *[Likert 1–5]*

Q14. Workforce skills and training (shortage of qualified AI talent, insufficient training programs) *[Likert 1–5]*

Q15. Risk management concerns (security vulnerabilities, liability uncertainties, unpredictable agent behavior) *[Likert 1–5]*

Q16. Model accuracy and reliability (hallucinations, inconsistent outputs, insufficient reasoning capabilities) *[Likert 1–5]*

Q17. Integration complexity (difficulty integrating with existing systems, APIs, and workflows) *[Likert 1–5]*

Q18. Cost constraints (high compute costs, unclear ROI, budget limitations) *[Likert 1–5]*

Q19. Regulatory compliance (existing or anticipated regulations, industry-specific requirements) *[Likert 1–5]*

Section 4: Governance Readiness and Outlook (Q20–Q22)

Q20. What is the current status of AI governance at your organization? *[Single choice: Formal AI governance framework in place / Partial framework (some policies but not comprehensive) / No formal governance framework]*

Q21. How does your organization measure the return on investment (ROI) of agentic AI initiatives? *[Single choice: Formal ROI tracking framework / Informal or anecdotal assessment / No ROI tracking]*

Q22. How soon do you expect your organization to reach L3 (Consultative) maturity or above? (If already at L3+, select ‘Already achieved’) *[Single choice: Already achieved / Within 1 year / 1–2 years / 2–3 years / Unlikely within 3 years]*

APPENDIX B. DETAILED SAMPLE CHARACTERISTICS

This appendix provides the full demographic profile of the survey sample (N = 136). All respondents met the inclusion criteria: direct involvement in AI strategy or implementation at organizations with 50 or more employees.

TABLE III. DISTRIBUTION BY INDUSTRY SECTOR (N = 136)

Industry Sector	n	%
Technology	37	27.2
Financial Services	26	19.1
Manufacturing	19	14.0
Retail & E-commerce	15	11.0
Healthcare & Pharmaceuticals	7	5.1
Telecommunications	6	4.4
Energy & Utilities	5	3.7
Education	5	3.7
Government & Public Sector	4	2.9
Professional Services & Consulting	4	2.9
Transportation & Logistics	3	2.2
Media & Entertainment	2	1.5
Real Estate & Construction	2	1.5
Other	1	0.7
Total	**136**	**100.0**

TABLE IV. DISTRIBUTION BY RESPONDENT ROLE (N = 136)

Role	n	%
Product Manager	35	25.7
Engineering Lead / Team Lead	30	22.1
CTO / VP of Technology	23	16.9
AI/ML Engineer	22	16.2
Business Analyst	16	11.8
Data Scientist	5	3.7
Other (IT Director, Architect, etc.)	5	3.7
Total	**136**	**100.0**

TABLE V. DISTRIBUTION BY ORGANIZATION SIZE (N = 136)

Organization Size	n	%
< 500 employees	40	29.4
500–5,000 employees	46	33.8
> 5,000 employees	50	36.8
Total	**136**	**100.0**

TABLE VI. GEOGRAPHIC DISTRIBUTION (N = 136)

Country	n	%
Russia	48	35.3
United States	20	14.7
United Kingdom	14	10.3
Germany	11	8.1
India	11	8.1
United Arab Emirates	8	5.9
Netherlands	7	5.1
Singapore	7	5.1
Brazil	10	7.4
Total	**136**	**100.0**

TABLE VII. SELF-ASSESSED EAAMM MATURITY LEVEL (N = 136)

EAAMM Level	n	%
L1 — Assistive	56	41.2
L2 — Collaborative	43	31.6
L3 — Consultative	26	19.1
L4 — Supervisory	8	5.9
L5 — Autonomous	3	2.2
Total	**136**	**100.0**

TABLE VIII. AI GOVERNANCE FRAMEWORK STATUS (N = 136)

Governance Status	n	%
Formal AI governance framework	24	17.6
Partial framework	47	34.6
No framework	65	47.8
Total	**136**	**100.0**