

From Meter-to-Cash to Agent-to-Agent Utilities: An Enterprise Architecture Framework for Autonomous Revenue Management in Sap-Based Utility Operations

Vikas Reddy Vallakonda

Independent Researcher, USA

ARTICLE INFO

ABSTRACT

Utility revenue-cycle systems built on SAP Meter-to-Cash architectures, spanning Customer Information Systems, Meter Data Management, and Financial Contract Accounting (FI-CA), were designed for periodic, deterministic billing cycles. The proliferation of distributed energy resources, electric vehicles, and large flexible loads such as artificial intelligence data centers has introduced revenue exposure that shifts faster than these systems were built to govern. This paper develops an enterprise architecture framework for delegating specific revenue-management functions, billing exception handling, revenue protection, settlement, and credit risk assessment, to autonomous software agents operating within a governed SAP execution layer (FI-CA, Business AI, Joule, Agent Builder). Rather than treating the entire utility enterprise as a single agent ecosystem, the framework isolates the revenue-management layer and specifies a tiered decision-authority model in which agent autonomy varies by decision type: full authority for reversible, high-volume decisions; shared authority for financially material actions; and no autonomous authority for disconnections, rate filings, or regulatory submissions. The framework is examined through an illustrative case constructed for a representative mid-sized North American utility and grounded in enterprise implementation experience with SAP FI-CA exception remediation. Findings suggest that governed agent delegation can compress revenue-protection response times materially relative to manual escalation processes, although the illustrative figures presented are exploratory rather than empirically validated. The paper's contribution is a decision-authority model specific to utility revenue management, extending prior enterprise-architecture and agent-governance literature, which has largely treated utility digitalization as an undifferentiated whole rather than isolating the financial control layer.

Keywords: Enterprise architecture; Revenue management systems; Agent-to-agent AI; Meter-to-cash; SAP Business AI; Utility digital transformation

1. INTRODUCTION

Utility revenue-cycle management has, for most of the industry's history, run on a predictable cadence. A meter reads consumption, a bill calculates charges against tariff rules, and a financial contract-accounting ledger posts and collects the result. SAP's Meter-to-Cash architecture, spanning the Industry Solution for Utilities (IS-U), Convergent Invoicing, and Financial Contract Accounting (FI-CA), was engineered for exactly this cadence: periodic, rule-based, and largely deterministic.

That cadence no longer matches the underlying business. Distributed solar generation, behind-the-meter battery storage, electric vehicle charging, and hyperscale artificial intelligence data centers introduce revenue exposure that moves in minutes rather than billing cycles. Global electricity demand tied to AI-optimized data centers is projected to grow substantially over the next decade (International Energy Agency, 2025), and a single curtailment event, net-metering true-up, or demand-response settlement can carry more financial materiality than an entire month of routine billing activity, even though the systems that would need to price, settle, and reconcile such events remain organized around monthly and quarterly processing windows. This shift is part of a broader pattern in which digital technologies disrupt established value-creation paths and force organizational responses (Vial, 2019), though most digital-transformation research examines this disruption at the level of the whole organization rather than within a single financially sensitive subsystem such as revenue management.

Two literatures speak to this problem without fully resolving it for the revenue-management function specifically. Enterprise architecture research treats governance and system-of-record integrity as first-order design concerns (Zachman, 1987), but predates the practical emergence of autonomous software agents capable of executing, not merely recommending, financial decisions. Multi-agent systems research (Wooldridge & Jennings, 1995) and its more recent extension into agentic artificial intelligence establish that software agents can coordinate complex, distributed decisions, but this literature has been applied to utility operations largely at the level of the whole enterprise: grid balancing, asset maintenance, customer engagement, and regulatory compliance addressed together, with revenue management appearing as one function among several rather than as a governance problem in its own right. Enterprise AI adoption research reinforces that this gap has practical consequences: a recent global survey finds that 78 percent of organizations report using AI in at least one business function while only 39 percent report measurable enterprise-level financial impact, and 62 percent are experimenting with AI agents specifically without yet having scaled them (Singla et al., 2025), a pattern consistent with agent deployments that are broad in scope but shallow in the governance detail any single function requires.

This gap matters because revenue management carries distinct constraints that grid operations or asset management do not share to the same degree. Billing and settlement actions are financially and legally consequential, subject to state and federal reporting obligations, and directly visible to customers in ways that predictive maintenance scheduling, for instance, is not. An architecture appropriate for delegating asset-inspection scheduling to an agent is not automatically appropriate for delegating a customer disconnection or a rate settlement.

This paper addresses that gap by developing a decision-authority framework specific to the revenue-management layer of utility enterprise architecture. The framework builds on SAP's existing Meter-to-Cash components as the governed system of record, and specifies how autonomous agents, built on SAP Business AI, Joule, and Agent Builder, can be granted full, shared, or no decision authority depending on the financial and regulatory characteristics of the decision type. The framework's propositions are examined against an illustrative case constructed for a representative utility and grounded, where possible, in enterprise implementation experience with FI-CA exception remediation. The remainder of the paper reviews the relevant architecture and agent-governance literature, develops three propositions about authority allocation, describes the illustrative case method, presents the framework and case analysis, and closes with theoretical and managerial implications.

2. THEORETICAL BACKGROUND AND PROPOSITIONS

2.1. Enterprise architecture governance in regulated utilities

Enterprise architecture frameworks were developed to give large organizations a disciplined way to align business process, data, application, and technology layers (Zachman, 1987). The Open Group's TOGAF standard formalized this into a repeatable architecture development method, and utilities have historically applied it to sequence large systems investments such as IS-U and FI-CA rollouts (The Open Group, 2022). Broader digital-technology adoption in the energy sector, including AI, is already visible through labor-market signals such as technology-specific job postings (Lyu & Liu, 2021), suggesting utilities are absorbing AI capability faster than their architecture-governance practices have adapted to supervise it.

What TOGAF-style governance does not by itself address is decision authority once a system is live: architecture frameworks specify how systems should be built and integrated, not which decisions inside those systems may be executed without a human in the loop. As utilities layer AI agents onto existing FI-CA and IS-U environments, this gap becomes operationally relevant rather than purely theoretical, because the enterprise architecture that governs system design increasingly needs to also govern agent authority.

2.2. From robotic process automation to autonomous agents in revenue-cycle management

Robotic process automation was the first widely adopted mechanism for reducing manual effort in utility and financial back-office functions, and its adoption record in adjacent industries is informative. A qualitative study of 14 RPA leaders and a survey of 139 RPA frontline workers in Big 4 accountancy firms found that automation changes the character of work more than it eliminates it and that RPA leaders and staff differ on whether the impacts are perceived positively or negatively (Cooper et al., 2022). Other work has directly considered risks, explaining that RPA scripts automating an imperfect business process simply mimic the process more rapidly and advocating governance frameworks to prevent RPA scripts propagating flaws at scale (Eulerich et al., 2024).

Field studies support the findings of dedicated governance: interviews with 24 experts (across 19 interviews) from 16 large, US- and European-based companies show that organizations that put formal ownership, monitoring, and escalation rules around their automation estate have materially fewer downstream incidents than those for whom governance comes as an afterthought (Kokina et al., 2026). Meanwhile, although there has been meaningful

experimentation and research into multi-agent systems with similar utility functions such as demand-response aggregation (e.g. with a bibliometric analysis of 587 Scopus-indexed documents supporting 39 core studies), the same is noticeably not true for the revenue-management layer.

Agentic artificial intelligence, in which large language model-based agents plan and execute multi-step tasks rather than following a fixed script, extends this trajectory but does not remove the underlying governance requirement; recent syntheses of the agentic AI literature explicitly flag opaque decision-making and limited human oversight as unresolved risks rather than solved problems (Brohi et al., 2025), and broader systematic reviews of agentic intelligence research identify governance and oversight as the least mature dimension of the field relative to technical capability (Tirulo et al., 2026). Hammer's (1990) earlier observation that organizations should redesign a process before automating it, rather than simply accelerating a flawed one, applies with particular force to revenue management, where the underlying FI-CA exception-handling logic often predates the automation layer being placed on top of it.

2.3. SAP FI-CA/BRIM as a governed execution layer for financial agents

SAP's Financial Contract Accounting module, together with the Billing and Revenue Innovation Management (BRIM) components (Subscription Order Management, Convergent Charging, and Convergent Invoicing), already function as the system of record for utility revenue transactions: rate determination, invoice calculation, payment application, and collections. This existing structure matters for agent governance because autonomous revenue-management agents do not need to be granted direct write access to a ledger of record; they can instead be scoped to propose, validate, or execute specific transaction types within FI-CA's existing authorization and audit-trail structure, operating under the same information-security management controls the utility already maintains for its core financial systems (ISO/IEC, 2022), with SAP Business AI and Joule providing the reasoning layer and Agent Builder providing the orchestration layer.

This is a materially different governance posture from agent architectures built on grid-control systems, where the operational control layer is deliberately kept outside AI participation because of millisecond-scale reliability requirements. Revenue-management decisions typically tolerate latency measured in minutes rather than milliseconds, which is precisely the latency band in which agent-based reasoning is best suited to operate.

2.4. Propositions

Three propositions follow from this reasoning and organize the framework developed in Section 4.

P1: Decision authority for revenue-management agents should be allocated by decision type rather than by function, because reversibility and financial materiality vary more within revenue management than between it and other utility functions.

P2: Agent authority appropriate for a legacy, deterministic Meter-to-Cash process is not automatically appropriate once that process is extended to cover real-time DER, EV, or large-flexible-load revenue exposure, because the volume and speed of exceptions changes even when the underlying financial risk category does not.

P3: Governance mechanisms proven in adjacent back-office automation contexts, explicit ownership, audit logging, and staged authority escalation (Kokina et al., 2026), transfer to agentic revenue management, but the specific authority thresholds must be set by decision type rather than inherited wholesale from RPA-era governance models built for narrower automation.

3. RESEARCH APPROACH

This paper adopts a design-oriented conceptual method common in enterprise-architecture research when a novel system class precedes the availability of representative deployment data: a framework is constructed from architecture and agent-governance literature, refined against the constraints described by relevant standards and regulatory bodies (ISO/IEC, NERC), and then examined through a single illustrative case rather than a statistically generalizable sample. This approach differs from the hypothesis-testing, survey-based, and panel-data designs more typical of empirical innovation-management research; the trade-off is explicit and is treated here as a limitation rather than a hidden weakness (see Section 5.3).

The illustrative case constructed in Section 4.3 uses a representative utility profile, 1.5 million customer accounts and US\$2.5 billion in annual revenue, consistent in scale with mid-sized North American investor-owned utilities, and draws on two categories of figures. The first category is the author's own enterprise implementation experience with SAP FI-CA exception remediation, specifically a collective-bill-print out-of-balance remediation program in which the out-of-balance bill rate was reduced from approximately 10 percent to approximately 2 percent through redesigned exception-handling logic; this figure is reported as the author's primary implementation experience

rather than as a benchmark drawn from a published source. The second category is illustrative order-of-magnitude financial modeling, constructed by the author for exposition purposes to demonstrate how a tiered agent-authority model could plausibly affect revenue-protection response time and cost; these figures are not derived from a specific utility's financial statements and are labeled as illustrative throughout Section 4.3 and in the accompanying tables.

Table 1 summarizes the comparison method used to contrast the legacy Meter-to-Cash architecture with the proposed agent-governed alternative across five dimensions: decision latency, exception-handling method, audit mechanism, financial materiality threshold for human review, and primary system of record. This comparison structure, rather than a statistical model, is the paper's principal analytical instrument.

Table 1. Legacy Meter-to-Cash vs. Agent-Governed Revenue Management

Dimension	Legacy Meter-to-Cash	Agent-Governed Revenue Management
Decision latency	Batch/periodic (daily to monthly cycles)	Near-real-time (minutes)
Exception handling	Manual queue review by billing staff	Agent-triaged, tiered by financial materiality
Audit mechanism	Transaction log within FI-CA	FI-CA audit trail plus agent decision log (Joule/Agent Builder)
Human-review threshold	Applied uniformly regardless of materiality	Applied selectively per the authority tiers in Table 2
Primary system of record	SAP IS-U / FI-CA	SAP IS-U / FI-CA (unchanged); agents operate as a reasoning layer above it

Note: comparison constructed by the author from the architecture and governance literature reviewed in Section 2.

The propositions developed in Section 2.4 are not tested statistically; they are used, in the manner common to conceptual and design-science contributions, as organizing claims that the framework and case analysis in Section 4 are built to satisfy and that future empirical work, outlined in Section 5.3, would need to test directly.

4. FRAMEWORK AND CASE ANALYSIS

4.1. Revenue Management Agent architecture

The framework specifies a Revenue Management Agent layer positioned above SAP FI-CA and BRIM, organized into four responsibility clusters. Billing exception handling covers out-of-balance bill detection, collective-bill-print reconciliation, and move-in/move-out proration errors, functions that, in the author's implementation experience, remain among the highest-volume sources of manual rework in legacy IS-U/FI-CA environments. Revenue protection covers theft detection, tamper flagging, and billing-anomaly identification, extending the anomaly-detection logic already established in adjacent financial-services applications of explainable machine learning, demonstrated in one representative study across an empirical analysis of 15,000 small and medium enterprise borrowers (Bussmann et al., 2021), a pattern consistent with the broader adoption of AI across financial services more generally (Vuković et al., 2025). Settlement and dynamic pricing covers demand-response incentive calculation and market-settlement exposure, the function most directly affected by DER, EV, and large-flexible-load growth. Credit risk assessment covers delinquency scoring and security-deposit determination, again drawing on methods established in the credit-risk literature rather than utility-specific precedent, since utilities have historically under-invested in this function relative to consumer lending (Bussmann et al., 2021).

Each cluster is instrumented through SAP Joule as the natural-language reasoning interface, SAP Business AI as the underlying predictive and anomaly-detection models, and SAP Agent Builder as the orchestration layer that sequences agent actions and writes results back into FI-CA through existing authorization objects. This positioning

keeps FI-CA as the authoritative ledger and audit trail, consistent with the governance posture described in Section 2.3, while allowing the reasoning and orchestration layers to be extended, retrained, or replaced independently of the underlying financial system of record, an important property given how quickly the underlying AI models are likely to change relative to the multi-year replacement cycle of a utility's core billing system.

4.2. Decision-authority boundaries

Not every revenue-management decision carries the same financial or regulatory weight, and the framework's central claim (P1) is that authority should track this variation directly rather than being set uniformly for the function as a whole. Table 2 classifies nine representative decision types into three authority tiers.

Table 2. Decision Authority Matrix for Revenue-Management Agent Actions

Decision Type	Agent Authority	Rationale
Standard billing inquiry response	Full	High volume, fully reversible, no financial exposure
Routine payment arrangement offer	Full	Reversible, governed by pre-approved terms
Demand-response incentive calculation (within approved bands)	Full	Bounded financial exposure, pre-approved pricing logic
Billing / out-of-balance exception remediation	Full	Reversible correction within existing FI-CA rules
Settlement posting above materiality threshold	Shared	Financially material; requires finance officer approval
Credit-limit / security-deposit adjustment	Shared	Affects customer financial obligation; moderate reversibility
Revenue-protection case escalation to field investigation	Shared	Triggers cost and potential legal exposure
Customer disconnection	None	Irreversible, legally regulated action
Rate case filing / regulatory submission	None	Legally consequential, regulator-facing

Note: authority tiers reflect the framework developed in Section 4.2; “Shared” denotes agent-proposed, human-approved.

Full agent authority is appropriate for high-volume, reversible, low-materiality decisions: routine payment-arrangement offers, standard billing-inquiry responses, and demand-response incentive calculations within pre-approved pricing bands. Shared authority, agent-proposed and human-approved, is appropriate where financial materiality is higher but the decision remains routine in structure: settlement postings above a threshold, credit-limit adjustments, and revenue-protection case escalation to field investigation. No autonomous authority is retained for customer disconnections, rate case filings, and regulatory submissions, consistent both with the general governance principle that irreversible or legally consequential actions remain under human authority regardless of how mature the underlying model becomes (Tabassi, 2023) and with critical-infrastructure protection standards that keep disconnection and emergency actions under direct human authority in North American utility markets (North American Electric Reliability Corporation, 2024).

This tiering is deliberately narrower in scope than utility-wide agent-governance proposals, which typically classify decisions at the level of the whole enterprise; scoping the matrix specifically to revenue management makes the

thresholds more precise and more defensible to a utility's finance and compliance functions, which are the primary audience for this particular authority model.

4.3. Illustrative case: revenue exposure management during a demand-response event

To examine how the framework's authority tiers behave under a representative event, consider a demand-response settlement event at a hypothetical mid-sized ERCOT-region utility with 1.5 million customer accounts and US\$2.5 billion in annual revenue, a scale consistent with typical investor-owned utility profiles in this market. Under the traditional process, a finance team member identifies a settlement exposure, manually calculates incentive pricing, contacts affected large customers, and negotiates participation; this sequence, drawn from typical utility demand-response operating practice, has been reported to take four to eight hours end to end with roughly 10 MW of curtailment typically achieved.

Under the proposed agent-governed alternative, a Revenue Management Agent operating within the full-authority tier calculates incentive pricing and settlement exposure automatically once a triggering condition is detected, a shared-authority step routes the resulting offer to a finance officer for approval given its financial materiality, and the offer is transmitted to the customer's own energy-management system. In this illustrative scenario, response time compresses to approximately 20 minutes and curtailment achieved increases to approximately 30 MW, not because the agent is more accurate than a human analyst, but because the shared-authority approval step removes the multi-hour scheduling delay inherent in a manual escalation chain.

Table 3 extends this illustration into an annual financial model for the utility profile described above, isolating four revenue-management-specific benefit categories: revenue protection, billing-exception remediation, settlement and credit-risk optimization, and regulatory-reporting automation. These figures are the author's own illustrative construction for exposition purposes, informed qualitatively by the billing-exception remediation experience described in Section 3, and are not drawn from a specific utility's audited financial statements.

Table 3. Illustrative Annual Benefit Model for Revenue-Management Agent Adoption

Benefit Category	Illustrative Annual Value	Basis
Revenue protection (theft / billing-anomaly detection)	\$7 million	Illustrative framework assumption
Billing-exception remediation	\$5 million	Illustrative; qualitatively consistent with the author's FI-CA out-of-balance remediation experience ($\approx 10\% \rightarrow \approx 2\%$)
Settlement and credit-risk optimization	\$4 million	Illustrative framework assumption
Regulatory / financial reporting automation	\$2.5 million	Illustrative framework assumption
Total illustrative annual benefit	\$18.5 million	Sum of above; representative utility profile (1.5 million customers, \$2.5 billion annual revenue)

Note: all figures are illustrative order-of-magnitude assumptions constructed by the author for expository purposes and are not derived from a specific utility's audited financial statements.

The billing-exception remediation figure in Table 3 is qualitatively consistent with the author's own implementation experience reducing an out-of-balance bill rate from approximately 10 percent to approximately 2 percent through redesigned FI-CA exception logic, though the dollar figure itself is illustrative rather than a direct restatement of that program's financials.

5. DISCUSSION AND CONCLUSIONS

5.1. Theoretical contributions

This paper's principal theoretical contribution is narrowing the unit of analysis in agent-governance research from the utility enterprise as a whole to the revenue-management layer specifically. Prior conceptual work on agent-to-agent utility ecosystems treats finance and revenue functions as one node among several, alongside grid operations, asset management, and regulatory compliance, without examining whether the authority thresholds appropriate for, say, predictive maintenance scheduling (Scaife, 2024) are also appropriate for settlement and disconnection decisions.

The decision-authority matrix developed here (Table 2) suggests they are not: revenue-management decisions cluster more tightly around financial materiality and reversibility as governing variables than around the functional category of finance itself, which supports Proposition 1's claim that authority should be allocated by decision type rather than by function. This also extends enterprise-architecture theory (Zachman, 1987) in a specific direction: where classical enterprise-architecture frameworks separate business, data, application, and technology layers, the framework developed here suggests that agent-era architectures may need an explicit decision-authority layer that cuts across the traditional four, since authority boundaries do not map cleanly onto any single one of them.

5.2. Managerial implications

For utility finance and IT leaders, the framework's practical implication is that agent adoption in revenue management does not need to wait for utility-wide digital transformation to mature. Because FI-CA already functions as a governed system of record (Section 2.3), a narrower, revenue-management-specific agent deployment can proceed on a shorter timeline than the broader multi-year roadmaps typically associated with full utility digital transformation.

Table 4 sketches a four-stage adoption path consistent with this narrower scope: an advisory stage in which agents surface exceptions for human action; a full-authority stage limited to the lowest-materiality decision types identified in Table 2; a shared-authority stage extending to settlement and credit decisions; and a periodic re-scoping stage in which authority thresholds are revisited as model performance and organizational trust accumulate, consistent with the general principle that autonomy should be earned incrementally rather than assumed (Tabassi, 2023).

Table 4. Phased Adoption Roadmap for Revenue-Management Agent Authority

Stage	Decision Types Included	Governance Mechanism
1. Advisory	All nine decision types (Table 2)	Agent surfaces a recommendation; a human executes every action
2. Full authority (low materiality)	Table 2, rows 1–4	Agent executes directly; sampled post-hoc audit review
3. Shared authority (moderate materiality)	Table 2, rows 5–7	Agent proposes; finance officer approves before execution
4. Periodic re-scoping	All tiers reassessed	Authority thresholds revisited on a fixed cycle as model performance and organizational trust accumulate

This staged approach also addresses a practical objection finance leaders are likely to raise: that agent errors in billing and settlement are more visible to customers and regulators than errors in, for instance, asset-maintenance scheduling. Beginning with the lowest-materiality, most reversible decision types allows a utility to build the audit-log and monitoring infrastructure that RPA governance research identifies as the strongest predictor of successful automation outcomes (Kokina et al., 2026) before extending authority to higher-materiality decisions.

5.3. Limitations and future research

This paper's principal limitation is methodological: the framework is examined through a single illustrative case rather than a representative sample, and the financial figures in Table 3 are explicitly exploratory rather than

empirically validated, a choice made deliberately to avoid presenting invented precision as if it were measured fact. This limits the claims that can be made about the magnitude of benefit from agent adoption, though it does not limit the paper's narrower claim about how authority should be allocated once adoption occurs.

A second limitation is scope: the framework addresses revenue management specifically and does not resolve authority allocation for grid operations or asset management, where different latency constraints apply. Future research could test Table 2 against actual deployment data as agentic revenue-management systems mature past the pilot stage, and could examine whether these thresholds hold across regulatory environments outside North America.

A further avenue, consistent with broader adoption-challenge literature in the energy AI space (Park, 2025), would be examining how utility regulators respond to agent-executed settlement decisions once they begin appearing in rate-case evidence, since regulatory acceptance, rather than technical capability, is likely to be the binding constraint on how far the shared-authority tier in Table 2 can be extended over time.

Ethical statement

This study did not involve human participants, animal subjects, or identifiable personal data requiring ethics board review. It presents a conceptual enterprise architecture framework and an illustrative case analysis based on secondary published sources and the author's own professional implementation experience, reported with appropriate attribution.

Data availability statement

No new empirical dataset was generated for this study. The illustrative financial figures presented in Section 4.3 and Table 3 are the author's own conceptual construction for expository purposes and are not derived from a proprietary or third-party dataset. All cited sources are listed in the reference list.

Declaration of competing interest

The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of Generative AI in Scientific Writing

During the preparation of this work, the author used AI-assisted tools to support literature search, drafting, and language editing. The author reviewed and edited all content and takes full responsibility for the content of this publication.

REFERENCES:

- [1] Brohi, S., Mastoi, Q., Jhanjhi, N.Z., & Pillai, T.R. (2025). A research landscape of agentic AI and large language models: Applications, challenges and future directions. *Algorithms*, 18(8), 499. <https://doi.org/10.3390/a18080499>
- [2] Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203–216. <https://doi.org/10.1007/s10614-020-10042-0>
- [3] Cooper, L.A., Holderness, D.K., Sorensen, T.L., & Wood, D.A. (2022). Perceptions of robotic process automation in Big 4 public accounting firms: Do firm leaders and lower-level employees agree? *Journal of Emerging Technologies in Accounting*, 19(1), 33–51. <https://doi.org/10.2308/JETA-2020-085>
- [4] Eulerich, M., Waddoups, N., Wagener, M., & Wood, D.A. (2024). The dark side of robotic process automation (RPA): Understanding risks and challenges with RPA. *Accounting Horizons*, 38(2), 143–152. <https://doi.org/10.2308/HORIZONS-2022-019>
- [5] Hammer, M. (1990). Reengineering work: Don't automate, obliterate. *Harvard Business Review*, 68(4), 104–112.
- [6] International Energy Agency. (2025). *Energy and AI*. IEA, Paris.
- [7] International Organization for Standardization. (2022). *ISO/IEC 27001:2022 — Information security management systems — Requirements*. ISO/IEC.
- [8] Kokina, J., Langmann, C., Kogan, G., & Weiser, M. (2026). Navigating the RPA governance landscape: Evidence from the field. *International Journal of Accounting Information Systems*, 57, 100784. <https://doi.org/10.1016/j.accinf.2026.100784>
- [9] Lyu, W., & Liu, J. (2021). Artificial intelligence and emerging digital technologies in the energy sector. *Applied Energy*, 303, 117615. <https://doi.org/10.1016/j.apenergy.2021.117615>
- [10] North American Electric Reliability Corporation. (2024). *Critical Infrastructure Protection (CIP) reliability standards*. NERC.

- [11] Olorunfemi, B.O., & Nwulu, N. (2023). Multi-agent system implementation in demand response: A literature review and bibliometric evaluation. *AIMS Energy*, 11(6), 1179–1210. <https://doi.org/10.3934/energy.2023054>
- [12] Park, C. (2025). Addressing challenges for the effective adoption of artificial intelligence in the energy sector. *Sustainability*, 17(13), 5764. <https://doi.org/10.3390/su17135764>
- [13] Scaife, A.D. (2024). Improve predictive maintenance through the application of artificial intelligence: A systematic review. *Results in Engineering*, 21, 101645. <https://doi.org/10.1016/j.rineng.2023.101645>
- [14] Singla, A., Sukharevsky, A., Kerr, D., Hjartar, K., Hämäläinen, L., Bout, S., Di Leo, V., & Dagorret, G. (2025). The state of AI: How organizations are rewiring to capture value. McKinsey & Company, QuantumBlack.
- [15] Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
- [16] The Open Group. (2022). The TOGAF Standard, 10th Edition. The Open Group.
- [17] Tirulo, A., Yadav, M., Lolamo, M., Chauhan, S., Siano, P., & Shafie-khah, M. (2026). Beyond automation: Unveiling the potential of agentic intelligence. *Renewable and Sustainable Energy Reviews*, 226, 116218. <https://doi.org/10.1016/j.rser.2025.116218>
- [18] Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *The Journal of Strategic Information Systems*, 28(2), 118–144. <https://doi.org/10.1016/j.jsis.2019.01.003>
- [19] Vuković, D.B., Dekpo-Adza, S., & Matović, S. (2025). AI integration in financial services: A systematic review of trends and regulatory challenges. *Humanities and Social Sciences Communications*, 12(1). <https://doi.org/10.1057/s41599-025-04850-8>
- [20] Woltmann, S., & Kittel, J. (2022). Development and implementation of multi-agent systems for demand response aggregators in an industrial context. *Applied Energy*, 314, 118841. <https://doi.org/10.1016/j.apenergy.2022.118841>
- [21] Wooldridge, M., & Jennings, N.R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. <https://doi.org/10.1017/S0269888900008122>
- [22] Zachman, J.A. (1987). A framework for information systems architecture. *IBM Systems Journal*, 26(3), 276–292. <https://doi.org/10.1147/sj.263.0276>