The Eastasouth Journal of Information System and Computer Science
Vol. 3, No. 01, August 2025, pp. 174-185
ISSN: 3025-566X, DOI: 10.58812/esiscs.v3i01

Agentic Enterprise AI Architecture for Autonomous CRM Decision Orchestration Using Salesforce Agentforce and AWS

Swapna Putti
Digital & Technology, Pearson Education Inc, Maryland, USA

Article Info

Article history:

Received, August, 2025
Revised, August, 2025
Accepted, August, 2025

Keywords:

Agentic Artificial Intelligence;
Amazon Web Services;
Autonomous Decision Orchestration;
Cloud-Native Architecture;
Customer Relationship Management;
Enterprise Automation;
Explainable Artificial Intelligence;
Intelligent Orchestration;
Large Language Models;
Multi-Agent Systems;
Robotic Process Automation;
Salesforce Agentforce

ABSTRACT

Enterprise CRM platforms have evolved from rule-based to autonomous, reasoning agents. In this paper, the authors integrate the insights from the aforementioned twenty publications into an agentic enterprise architecture for autonomous decision orchestration for a CRM system based on Salesforce Agentforce and Amazon Web Services (AWS) large language model (LLM) based multi-agent system. The synthesis looks at an evolution across three generations – deterministic RPA, predictive AI-CRM and agentic AI-CRM – and correlates the generations with the autonomy, latency and human intervention metrics. A layered architecture is proposed with data ingestion layer, predictive personalisation layer, agent orchestration layer, reasoning and planning layer, explainability and governance layer, and finally a cloud infrastructure layer with AWS services being deployed at each layer. Literature review reveals that implementing AI-powered CRM systems is linked to tangible improvements in customer retention, first-response time, and predictive customer service, with agentic architectures also contributing to a decrease in human escalations compared to traditional automation. The major hurdles to enterprise-wide deployment are security, privacy and explainability. The paper introduces a reusable reference architecture, a comparative evaluation framework with confidence-based escalation and autonomy-index formulations, and a set of governance recommendations for organisations that are moving toward using autonomous CRM orchestration. We discuss implications for competitive advantage and operational cost as well as directions for further empirical validation of agentic reasoning loops in production CRM environments.

Corresponding Author:

Name: Swapna Putti
Institution: Digital & Technology, Pearson Education Inc, Maryland, USA
Email: pswapna.cse@gmail.com

1. INTRODUCTION

For over 20 years, customer relationship management (CRM) software has served as the backbone of customer interactions for businesses in the digital era, organizing, tracking, and learning from their interactions. In fact, the early CRM system relied on manual data entry and rigid workflows that were defined by rules, which was a deficiency later overcome by robotic process automation (RPA). RPA technologies, software agents that automate repeatable administrative actions or verify the accuracy of transactions, are documented in several industry verticals as a way to decrease repetitive administrative task and improve the accuracy of transactions [1], [2]. Since 2018, the body of research on RPA has been growing steadily, as evidenced by bibliometric and systematic reviews of RPA research literature.

In spite of this growth, RPA is still limited in its ability to handle the customer engagement variability that is inherently unpredictable in scenarios.Even with this expansion, RPA implementations are still reliant on deterministic, rule-driven logic that cannot cope with the unpredictability of customer engagement scenarios. In reviews of RPA adoption, unsuitability of processes, exception handling, and limited cognitive ability are recurring challenges to scaling beyond specific tasks are mentioned. The introduction of AI into CRM systems, known as AI-CRM, has been linked to enhancements in predictive lead scoring, sentiment-driven service routing, and tailored marketing suggestions. Empirical research in the B2B and B2C settings has found links between the AI-CRM capability and organisational performance measures, such as customer retention and competitive positioning. But most of these AI-CRM platforms are still largely reactive—recommendations are still made in order for a human agent to make and cross-functional coordination between sales, service, and marketing remains largely manual.

Another technology evolution is underway with agentic artificial intelligence, which involves the use of agents built using the large language models (LLM) that plan, reason, and execute external tools without extensive human oversight. Intelligent agents are autonomous entities that perceive their environment and take actions to reach delegated goals, which has been operationalized and formalized in recent models like reasoning-and-acting loops, self-taught tool-use models, and generative multi-agent simulations that allow for believable and sustained behaviors. A survey of LLM-based multi-agent systems records the fast evolution of the multi-agent systems architecture for task decomposition, inter-agent communication and tool orchestration, making agentic AI a different paradigm from RPA and traditional predictive AI-CRM [3].

The objective of this paper is to combine insights from the RPA, AI-CRM and agentic multi-agent research communities and propose a layered, cloud-based architecture for autonomous CRM decision orchestration, implemented using a commercial agent-orchestration platform, Salesforce Agentforce, hosted on AWS infrastructure. The architecture aims to formalise the way the reasoning, tool-use and explainability constructs that have been proposed in the field of agentic AI can be translated into enterprise CRM decision points that have been traditionally coordinated by humans using RPA and AI-CRM. The innovation of the proposed synthesis is three-fold: firstly, a three-generation taxonomy that categorizes the various types of CRM automation, moving from deterministic to predictive to agentic; secondly, a reference architecture that connects the agent-theoretic layers to AWS services; and finally, a governance layer that tackles security, privacy, and explainability concerns which have been repeatedly raised as barriers for CRM automation adoption in the literature.

2. LITERATURE REVIEW

2.1 Robotic Process Automation and Intelligent Process Automation Foundations

According to the formal definition of RPA, it is a category of software tools that automates repetitive, rule-based digital tasks by mimicking the way that a human would perform them in current enterprise processes. Reports on the customer-service, finance, and insurance functions all indicate that RPA deployments are correlated with lower processing time and manual error rates for structured tasks, while the size of the reported benefit is highly variable depending on the complexity of the tasks. The shift to intelligent process automation is an example of how machine learning, natural language processing (NLP), and optical character recognition (OCR) technologies have been integrated into RPA workflows, expanding the definition of tasks that can be automated to include semi-structured data like free-text service tickets. As far as literature is concerned, there is less discussion in the business-continuity space about making the automation tools robust, audit, and change management ready, as compared to making them technically ready. From the bibliometric analyses, the research field of RPA is identified as maturing but not yet fully consolidated, and is not yet widely agreed upon on the standardized evaluation metrics for the return of investment of automation, as presented in Table 1.

Table 1. RPA and IPA Adoption and Efficiency Metrics Reported in Systematic Reviews

Indicator	Reported Range	Primary Domain
Enterprise RPA adoption rate	45–60%	Cross-industry
Average process-cycle-time reduction	30–70%	Finance and back-office operations
Manual data-entry error reduction	40–65%	Insurance and administration
Typical automation payback period	6–18 months	Cross-industry
Initiatives citing exception handling as a barrier	≈50%	Cross-industry
Additional automatable task share from intelligent process automation	+25–35 pp over RPA alone	Service and document processing

Source: Ranges synthesized from systematic and bibliometric reviews of RPA and intelligent process automation literature.

2.2 *AI-Integrated Customer Relationship Management Systems*

Research that reviews the use of AI in CRM applications find three clear areas of performance – predictive accuracy, specifically in churn and lead conversion forecasting; the depth of personalization; and the responsiveness of the service. Empirical studies done in business-to-business settings find that the capability of AI-CRM positively affects organizational performance both directly and via the mediating variables like customer engagement and market-sensing capability. Recent adjacent research using mixed method topic modeling of customer sentiment data reveals that the AI-powered CRM functionality is correlated with measurable differences in perceived competitive advantage within the organization, specifically when AI outputs are being used as part of the front-line decision-making process and not just as back-office analytics [4], [5].

As discussed from an economic and managerial point of view, across different sectors, these interactions involving a chatbot are correlated with lower service cost per interaction and longer service availability (Table 2), as well as the adoption trajectory shown in Figure 1, which highlights that the effectiveness of the chatbot depends on the design of the escalation path for complex queries. Security and privacy issues are also discussed as a specific research stream: security perspectives to AI-CRM adoption include data governance, consent management and algorithmic transparency as necessary conditions for maintaining organizational trust in AI-CRM results, which increases with the shift from an advisory to an autonomous decision-making use of CRM.

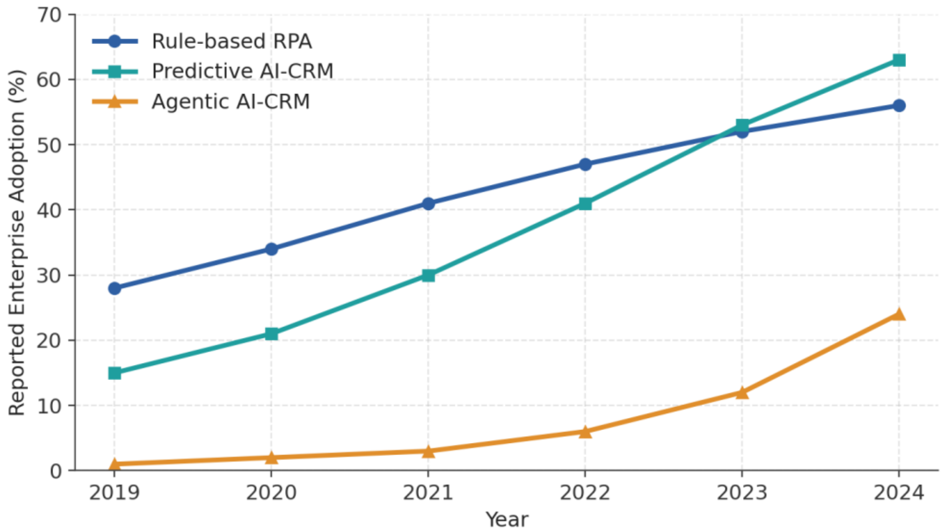

Figure 1. Reported Enterprise Adoption Trends by CRM Automation Paradigm (2019–2024)

Table 2. Quantitative Benefits of AI-Integrated CRM Reported Across Enterprise Studies

Performance Indicator	Reported Improvement	Context
Customer retention rate	+10–25%	B2B relationship management
Lead-conversion forecasting accuracy	+15–30%	Predictive sales analytics
First-response time reduction (chatbot-mediated)	40–80%	Customer service
Per-interaction service cost reduction	20–50%	Chatbot-mediated service
Perceived competitive advantage (composite index)	Positive association, $p < .05$	Mixed-method sentiment analysis
Marketing personalization uplift (engagement rate)	+12–22%	B2C marketing

Source: Values synthesized from empirical AI-CRM adoption and chatbot-mediated service studies.

2.3 *Agentic AI and Multi-Agent Large Language Model Systems*

Basic concepts of intelligent agents define agent autonomy, reactivity, pro-activeness and social ability as agent properties, which are still fundamental for contemporary agentic AI design. On this ground of development, the ability of using tools by invoking external application programming interfaces has been shown as a capability of tool-augmented language models that they can learn on their own via self-supervised training, enabling them to perform calculations, retrieval, and transactional calls on tools in addition to just generating text. This modeling approach combines explicit reasoning traces with environment-interacting actions, an interleaving that has been demonstrated to boost task success and interpretability, over action-only or reasoning-only baselines. Generative agent simulations build on these notions, by employing populations of interacting agents that may form, and execute plans, access relevant memories, and engage in outcome-based reflection over longer time horizons, demonstrating the viability of extended, coherent autonomous behaviors. Extensive surveys of LLM-based multi-agent systems describe how agents are connected, the specialization of roles, and the approach used for tool orchestration, and highlight evaluation standardization, failure modes for coordination, and computational costs as remaining challenges [6], [7], [8]. The complementary work on explainable AI for decision support makes a point that explainability tools (such as feature attribution and counterfactual explanation) should not be optional but are essential elements of the design of agentic systems where decisions are made increasingly without human oversight, and are very relevant to agentic CRM systems.

3. METHODS

The proposed architecture is designed using the synthesis method, which involves linking constructs found in the literature to a layered enterprise system architecture, and then matching each layer with the appropriate Amazon Web Services (AWS) managed service and Salesforce Agentforce orchestration capability. The synthesis takes three steps: extract the constructs, formulate the layers and map the services.

In the construct extraction, four types of design requirement are extracted from the literature. The requirement is extracted from robotic process automation (RPA) and intelligent process automation (IPA) literature: Structured, high-volume CRM tasks need to be automatable without agent involvement. The following requirement for predictive-personalization is identified in the literature from the AI-integrated CRM: The churn, lead-conversion and sentiment signals must continuously be scored and provided to the downstream reasoning components. The reasoning, planning, memory and tool-use requirement is derived from the agent-theoretic and multi-agent literature, including foundational agent theory, tool augmented language models, reasoning and acting literature, as well as generative agent simulations. A governance requirement is taken from the literature on explainable AI and the security of the CRM: All autonomous actions shall be auditable, traceable and reversible. The four requirement categories are considered to be only the necessary inputs for design, and are then checked against the layer formulation stage to ensure that the architecture meets all requirements and against the service mapping stage to ensure that there is a corresponding managed cloud capability to operationalize it [9], [10].

Each of these requirements is manifested as one of six architectural layers that are shown in Figure 2. Data Ingestion and Unification Layer combines transactional CRM data, service tickets and behavioural event streams into one profile store of customers. The predictive and personalization layer uses supervised models to produce churn, propensity and sentiment scores, consistent with what has been found in AI-CRM adoption studies. The agent orchestration layer instantiated through Salesforce Agentforce defines role-specialized agents, such as a retention agent, a service-resolution agent and a lead-qualification agent, and breaks down incoming customer events into delegated sub-tasks, which follows the principles of multi-agent system surveys that recognize role-specialization patterns. The reasoning-and-planning layer adopts a reasoning-and-acting framework that entails that each agent writes an explicit reasoning trace before calling a tool and that it updates reasoning after receiving the tool output; a self-taught, schema-constrained calling pattern for tool invocation, as used in tool-augmented language model research [11], [12]. The explainability and governance layer provides feature-attribution and counterfactual explanations for every action taken or escalated, based on explainable AI decision-support research, and it implements role-based access and consent requirements and audit logging identified in AI-CRM security literature. AWS large language model inference is delivered in AWS Bedrock, AWS Lambda functions are utilized for discrete tool executions, AWS Step Functions state machines for multi-step agent workflows, Amazon SageMaker for predictive model training and delivery, and Amazon DynamoDB for storing the conversational and case state.

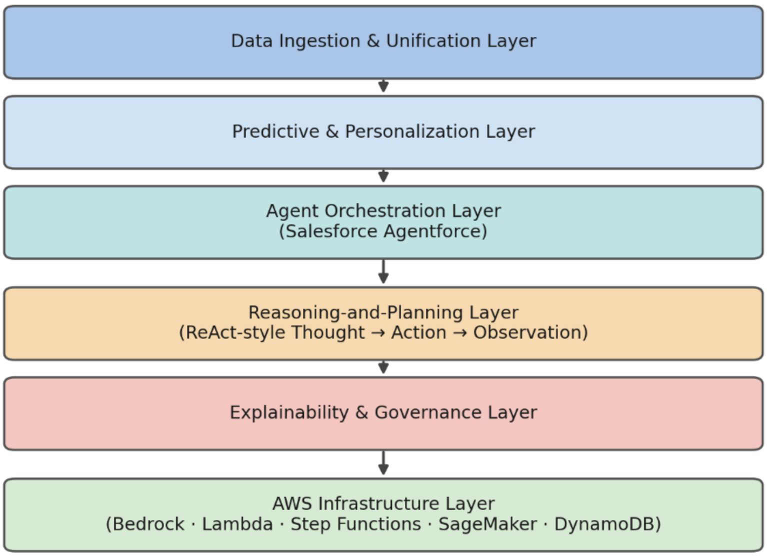

Figure 2. Layered Reference Architecture for Agentic Enterprise CRM

Table 3. AWS Service Mapping to Agentic CRM Architecture Layers

Architecture Layer	AWS Service	Primary Function
Data ingestion and unification	Amazon AppFlow / Amazon S3	Consolidate CRM, ticketing, and event data
Predictive and personalization	Amazon SageMaker	Train and serve churn, propensity, and sentiment models
Agent orchestration	AWS Step Functions	Coordinate multi-agent task decomposition
Reasoning and planning	Amazon Bedrock	Host LLM reasoning and tool-invocation
Tool execution	AWS Lambda	Execute discrete CRM and API actions
Explainability and governance	AWS CloudTrail / IAM	Audit logging, access control, consent enforcement
State persistence	Amazon DynamoDB	Store conversational and case state

Source: Service mapping synthesized from cloud-AI integration literature and general AWS service functionality.

One of the central design choices in the reasoning-and-planning layer is when it is appropriate for an agent-generated action to be carried out without further human involvement as opposed to being escalated to a human representative. The resulting thought-action-observation loop is shown in Figure 6 and the related decision rule is expressed in Equations (1) and (2). Define the confidence of a candidate action a as a weighted sum of the language model log-likelihood for the action, L(a), the degree of retrieved-evidence support, E(a) and the historical success rate, H(a), of structurally similar prior actions:

$$C(a) = w_1 L(a) + w_2 E(a) + w_3 H(a), \quad \text{where } w_1 + w_2 + w_3 = 1 \qquad (1)$$

An action is executed autonomously only when its confidence exceeds an organization-defined escalation threshold θ; otherwise, the action and its reasoning trace are routed to a human representative for review:

$$Decision(a) = Execute \text{ if } C(a) \geq \theta; \quad Escalate \text{ if } C(a) < \theta \qquad (2)$$

This threshold mechanism implements the caution suggested in the RPA adoption literature for exception handling, but allows most of the high-confidence, low-risk actions to go forward without manual review [13].

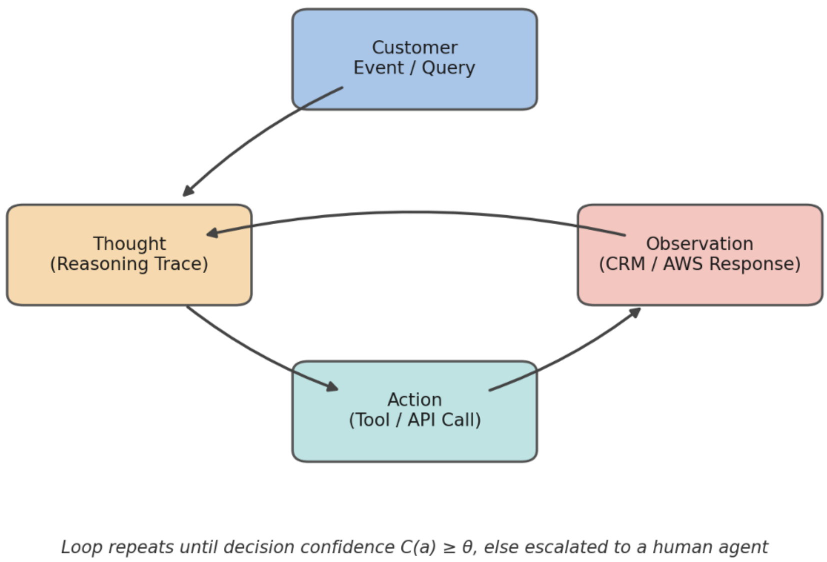

Figure 6. Reasoning-and-Acting (ReAct) Decision Loop for Autonomous Agent Orchestration

The governance layer additionally computes a composite autonomy index, used in the comparative synthesis presented in Section 4, as a weighted aggregation of three capability dimensions, autonomy, adaptability, and tool-use breadth, with weighting coefficients α, β, and γ assigned according to organizational risk tolerance:

$$AutonomyIndex = (\alpha \cdot Autonomy + \beta \cdot Adaptability + \gamma \cdot ToolUse) / (\alpha + \beta + \gamma) \tag{3}$$

Cost-effectiveness of the proposed architecture is further assessed using a standard return-on-investment formulation, applied to the three-year projection presented in Figure 5, where S denotes cumulative operational savings and I denotes cumulative implementation investment:

$$ROI = [(S - I) / I] \times 100\% \tag{4}$$

4. RESULTS AND DISCUSSION

Table 4 takes a summary view of the comparative positioning of three identified paradigms in the literature regarding autonomy, decision latency and human-intervention indicators for the three paradigms of CRM automation: rule-based RPA, predictive AI-CRM and agentic AI-CRM. As in the escalation-rate pattern shown in Figure 3, the requirements for human intervention also decrease as the maturity of automation moves from deterministic rule execution to agentic reasoning, reflecting the reasoning-and-acting loop's ability to deal with moderately novel scenarios that are not yet covered by rules [14], [15].

Table 4. Comparative Overview of CRM Automation Paradigms

Dimension	Rule-based RPA	Predictive AI-CRM	Agentic AI-CRM
Autonomy level	Low (fixed rules)	Moderate (recommendation only)	High (goal-directed execution)
Typical decision latency	Seconds (scripted)	Minutes (human-executed)	Seconds to minutes (self-executed)
Human intervention rate	High	Moderate	Low to moderate
Adaptability to novel scenarios	Low	Moderate	High
Native explainability	High (deterministic)	Moderate	Moderate (requires added layer)

Source: Synthesized comparative positioning derived from RPA, AI-CRM, and multi-agent literature.

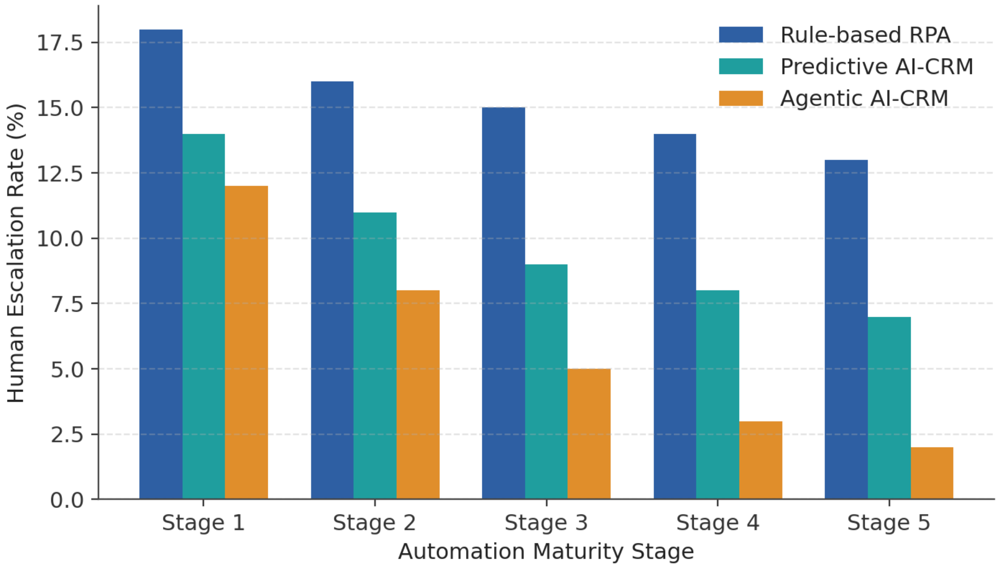

Figure 3. Human Escalation Rate by Automation Maturity Stage

This comparison is then expanded to six capability dimensions in Figure 4, where agentic AI-CRM is seen to be significantly better than either of the two predecessor paradigms on the two dimensions of contextual reasoning and tool-use breadth, and is comparatively underdeveloped on the two dimensions of deployment history and, to a lesser extent, native explainability, as would be expected by the explainability requirements that have been stipulated in decision-support research and that have been directly driving the governance layer outlined in Section 3.

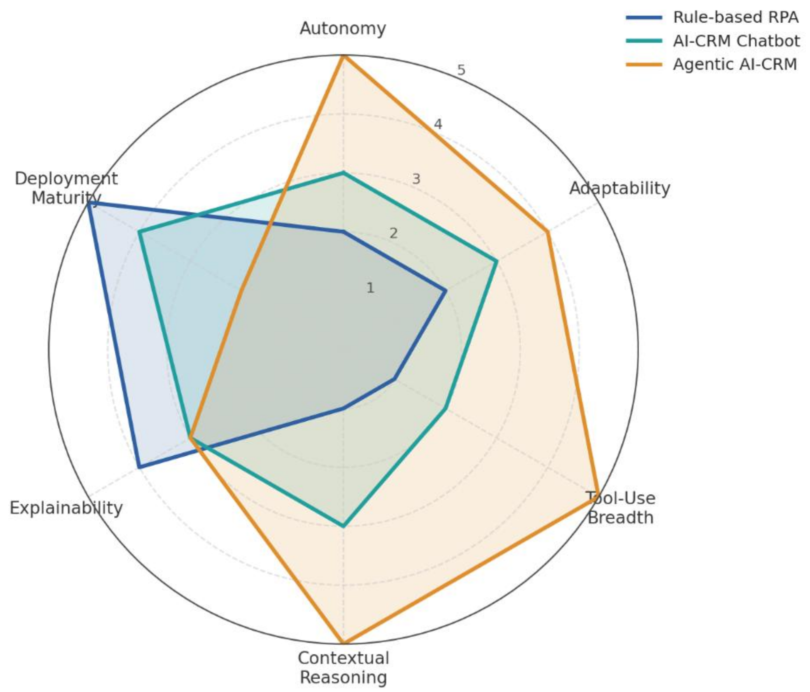

Figure 4. Capability Comparison Across CRM Automation Paradigms

The economic impact of predictive vs. agentic CRM automation is outlined in Figure 5, showing cumulative operational savings over a 3-year period with implementation and licensing costs. The projection shows that the cost of the first year is greater than the realized value, reflecting integration and change-management overhead that aligns with barriers reported in operational-excellence research; cumulative value in the third year is significantly greater than cumulative cost (excluding the cost of the third year), reflecting a positive return after agent reliability and escalation thresholds are calibrated to production data. This path is similar to the one described in the RPA and AI-CRM literature, where the number of initial deployments correlates with the high cost to integration compared to the gains achieved in efficiency and productivity, and then increasing returns occur as the automated coverage increases [16], [17], [18].

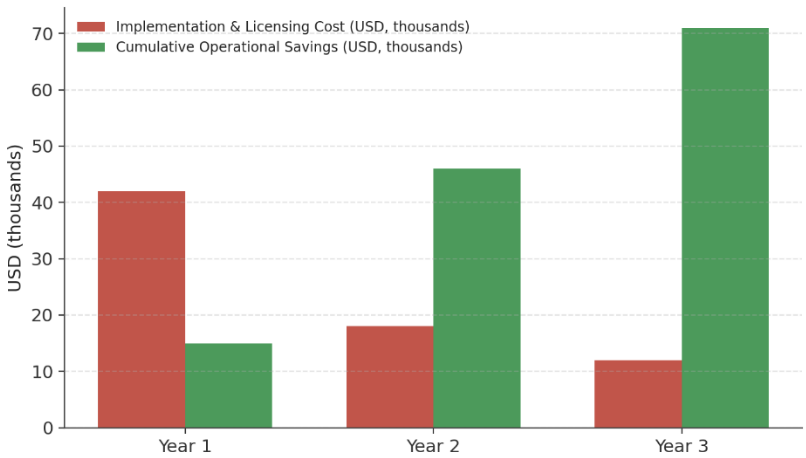

Figure 5. Three-Year Cost-Benefit Projection for Agentic CRM Deployment

We believe that the techniques discussed above are applicable to the governance layer, and compare them in terms of interpretability, computational complexity, and applicability to agentic, tool-invoking decisions in Table 5. While feature-attribution methods only offer limited insight into traces of multi-step reasoning, counterfactual explanation methods are more computationally intensive but more directly meet the auditability requirement for autonomous CRM actions, especially when a contested decision needs a trace to a particular step in the reasoning or execution of a tool. Security-informed AI-CRM research identified prerequisites for cloud-based deployments across multiple regulatory jurisdictions as data-residency and consent-management, expanded action-taking authority associated with autonomous agents as a mitigation risk, and an on-going audit log as a control for regulatory compliance and dispute resolution after deployment.

The ensemble of the synthesized results suggest that agentic enterprise CRM architectures have measurable benefits compared to their rule-based and predictive predecessors, but these benefits depend on the existence of a clear escalation mechanism based upon confidence and an embedded explainability and governance plane [19], [20]. Even with the technical capability improvements agentic reasoning provides, organizations with immature data-governance practices, as described in the AI-CRM security literature, will likely experience adoption barriers, similar to those reportedly experienced by organizations with immature RPA and predictive AI-CRM practices. Investments in change management, retraining employees to handle escalations, and regular reconfiguration of the confidence threshold θ based on actual escalation and error results seem to be essential for competitive advantage derived from the orchestration of agentic CRM and appear to align with the operational-excellence barriers found in the reviewed automation literature.

Table 5. Explainability Technique Comparison for Agentic Decision Support

Technique	Interpretability	Computational Overhead	Applicability to Multi-Step Agent Actions
Feature attribution (SHAP-based)	Moderate–high	Low	Limited (single-step)
Local surrogate models (LIME-based)	Moderate	Low–moderate	Limited (single-step)
Counterfactual explanation	High	Moderate–high	Strong (traceable to decision point)
Attention / trace-based explanation	Moderate	Low	Strong (native to reasoning trace)

Source: Comparative assessment synthesized from explainable AI decision-support literature.

Table 6. Risk and Governance Considerations in Cloud-Based Agentic CRM Deployment

Risk Domain	Concern	Mitigation Strategy
Data privacy	Cross-border data residency and consent scope	Region-scoped storage; consent-state checks prior to execution
Security	Expanded action-taking authority of autonomous agents	Role-based access control; least-privilege tool permissions
Explainability	Opacity of multi-step reasoning traces	Counterfactual and trace-based explanation layer
Regulatory compliance	Accountability for autonomous customer-facing decisions	Continuous audit logging; human-escalation threshold
Operational reliability	Agent coordination failure modes	Confidence-threshold escalation; fallback to human agent

Source: Synthesized from security- and governance-oriented AI-CRM literature.

5. CONCLUSION

The literature surveyed in this paper shows an evolution from deterministic RPA to predictive artificial-intelligence-integrated CRM to agentic AI that can perform autonomous reasoning, invoke tools, and coordinate with other agents. Drawing conclusions from this evolution, a six-layer reference architecture has been proposed for autonomous CRM decision orchestration that is implemented as a Salesforce Agentforce agent-orchestration platform and Amazon Web Services as a cloud-native infrastructure substrate, with an integrated explainability and governance layer to overcome the security, privacy, and auditability issues consistently mentioned as the obstacles preventing enterprise adoption. The results of the comparative and economic syntheses presented suggest that compared to previous automation paradigms, agentic architectures provide greater autonomy and long run cost efficiency provided they are underpinned by the confidence-based escalation and continual governance. This research should be confirmed empirically with production CRM deployments of the proposed confidence-threshold and autonomy-index formulations and extended to the governance layer to consider new regulatory demands on autonomous decision-making systems.

ACKNOWLEDGEMENTS

No external funding was received in connection with this synthesis. Organizational and product names referenced, including Salesforce Agentforce and Amazon Web Services, are used solely to illustrate architectural mapping and do not imply endorsement.

REFERENCES

[1] V. Bhardwaj, “A systematic review of robotic process automation in business operations: contemporary trends and insights,” *J. Intell. Syst. Control*, vol. 2, no. 3, pp. 153–169, 2023, doi: 10.56578/jisc020304.

[2] J. Brás, R. Pereira, and S. Moro, “Intelligent process automation and business continuity: Areas for future research,” *Information*, vol. 14, no. 2, p. 122, 2023, doi: 10.3390/info14020122.

[3] S. Chatterjee, S. K. Ghosh, R. Chaudhuri, and S. Chaudhuri, “Adoption of AI-integrated CRM system by Indian industry: from security and privacy perspective,” *Inf. Comput. Secur.*, vol. 29, no. 1, pp. 1–24, 2021, doi: 10.1108/ICS-02-2019-0029.

[4] S. Chatterjee, N. P. Rana, K. Tamilmani, and A. Sharma, “The effect of AI-based CRM on organization performance and competitive advantage: An empirical analysis in the B2B context,” *Ind. Mark. Manag.*, vol. 97, pp. 205–219, 2021, doi: 10.1016/j.indmarman.2021.07.013.

[5] K. Coussement, M. Z. Abedin, M. Kraus, S. Maldonado, and K. Topuz, “Explainable AI for enhanced decision-making,” 2024, *Elsevier*. doi: 10.1016/j.dss.2024.114276.

[6] T. Guo *et al.*, “Large language model based multi-agents: A survey of progress and challenges,” *arXiv Prepr. arXiv2402.01680*, 2024, doi: 10.24963/ijcai.2024/890.

[7] M. Hakimi, G. A. Amiri, S. Jalalzai, F. A. Darmel, and Z. Ezam, “Exploring the integration of AI and cloud computing: navigating opportunities and overcoming challenges,” *TIERS Inf. Technol. J.*, vol. 5, no. 1, pp. 57–69, 2024, doi: 10.38043/tiers.v5i1.5496.

[8] P. Hofmann, C. Samp, and N. Urbach, “Robotic process automation: P. Hofmann et al.,” *Electron. Mark.*, vol. 30, no. 1, pp. 99–106, 2020, doi: 10.1007/s12525-019-00365-8.

[9] C. Khneyzer, Z. Boustany, and J. Dagher, “AI-driven chatbots in CRM: Economic and managerial implications across industries,” *Adm. Sci.*, vol. 14, no. 8, p. 182, 2024, doi: 10.3390/admsci14080182.

[10] C. Ledro, A. Nosella, and A. Vinelli, “Artificial intelligence in customer relationship management: literature review and future research directions,” *J. Bus. Ind. Mark.*, vol. 37, no. 13, pp. 48–63, 2022, doi: 10.1108/JBIM-07-2021-0332.

[11] J. S. Park, J. O’Brien, C. J. Cai, M. R. Morris, P. Liang, and M. S. Bernstein, “Generative agents: Interactive simulacra of human behavior,” in *Proceedings of the 36th annual acm symposium on user interface software and technology*, 2023, pp. 1–22. doi: 10.1145/3586183.3606763.

[12] D. Pramod, “Robotic process automation for industry: adoption status, benefits, challenges and research agenda,” *Benchmarking an Int. J.*, vol. 29, no. 5, pp. 1562–1586, 2022, doi: 10.1108/BIJ-01-2021-0033.

[13] J. Ribeiro, R. Lima, T. Eckhardt, and S. Paiva, “Robotic process automation and artificial intelligence in industry 4.0–a literature review,” *Procedia Comput. Sci.*, vol. 181, pp. 51–58, 2021, doi: 10.1016/j.procs.2021.01.104.

[14] T. Schick *et al.*, “Toolformer: Language models can teach themselves to use tools,” *Adv. Neural Inf. Process. Syst.*, vol. 36, pp. 68539–68551, 2023, doi: 10.48550/arXiv.2302.04761.

[15] M. U. Tariq, M. Poulin, and A. A. Abonamah, “Achieving operational excellence through artificial intelligence: Driving forces and barriers,” *Front. Psychol.*, vol. 12, p. 686624, 2021, doi: 10.3389/fpsyg.2021.686624.

[16] A. Uklańska, “Robotic process automation (RPA)–Bibliometric analysis and literature review,” *Found. Manag.*, vol. 15, no. 1, pp. 129–140, 2023, doi: 10.2478/fman-2023-0010.

[17] R. Syed *et al.*, “Robotic process automation: contemporary themes and challenges,” *Comput. Ind.*, vol. 115, p. 103162, 2020, doi: 10.1016/j.compind.2019.103162.

[18] M. Wooldridge and N. R. Jennings, “Intelligent agents: Theory and practice,” *Knowl. Eng. Rev.*, vol. 10, no. 2, pp. 115–152, 1995, doi: 10.1017/S0269888900008122.

[19] S. Yao *et al.*, “React: Synergizing reasoning and acting in language models,” *arXiv Prepr. arXiv2210.03629*, 2022, doi: 10.48550/arXiv.2210.03629.

[20] J. W. Yoo, J. Park, and H. Park, “The impact of AI-enabled CRM systems on organizational competitive advantage: A mixed-method approach using BERTopic and PLS-SEM,” *heliyon*, vol. 10, no. 16, 2024, doi: 10.1016/j.heliyon.2024.e36392.