iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** XI **Month of publication:** November 2025

DOI: https://doi.org/10.22214/ijraset.2025.75852

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 13 Issue XI Nov 2025- Available at www.ijraset.com

Toward Agentic AI Frameworks for Technical Triage and Automated Product Support

Sri Krishna Ravulapalli
R&D Software Engineer, Broadcom, Texas, USA

Abstract: ***Agentic AI systems have achieved wide adoption in software development, assisting engineers with code generation, diagnostics, and efficiency. However, the evolution of AI-powered agents for support engineers and IT operations remains nascent. This paper surveys the landscape of agentic AI in both software development and operational support, examines current capabilities such as observability, root cause analysis, and explores the potential for automated remediation, knowledge-based case resolution, and SME-guided learning. We propose a taxonomy of support-centric AI agent tasks and discuss research and engineering challenges for autonomous, end-to-end product support by AI.***
Keywords: Agentic AI, Product support, IT operations, Observability, Knowledge Retrieval and Root Cause Analysis.

I. INTRODUCTION

Recent advances in artificial intelligence (AI) have transformed the roles of engineers in software development and product support. Agentic AI, defined as intelligent agents capable of autonomous decision-making and action, has seen rapid progress in coding assistants. However, analogous capabilities tailored for IT operations and support engineers are still evolving. This paper reviews the current state of agentic AI in an engineering workflow, highlights opportunities for operational and support automation, and outlines future research directions for comprehensive product support.

II. AGENTIC AI IN SOFTWARE DEVELOPMENT

AI-powered coding assistants are widely adopted by software developers. These agentic AI systems, trained in programming languages such as Java, Python, and NodeJs accept user requests in natural language and generate relevant code snippets. Integration of such agents accelerates development, improves code quality, and streamlines the completion of tasks for engineers.
Developers can leverage AI-generated code in the end-to-end system components, iteratively refining solutions, and reducing the cognitive burden of routine tasks. AI is providing developers with substantial benefits in numerous usage scenarios, such as:

1) Automated Planning and Requirement Analysis: Agentic AI accelerates project planning by automatically extracting requirements and dependencies from documentation and historical data. This reduces manual effort and increases accuracy in scoping and forecasting development timelines.
2) End-to-end Code Generation and Review: Beyond snippet generation, agentic AI autonomously builds complete modules, optimizes code logic, enforces coding standards, and performs static analysis to detect bugs and security vulnerabilities before human review. For example, an AI agent in JavaScript projects can identify inefficient code patterns and automatically replace them with optimized solutions.
3) Intelligent and Self-improving Testing: AI agents generate, run, and evolve test cases autonomously. They identify edge cases, detect regressions, and integrate into continuous integration/continuous deployment (CI/CD) pipelines for real-time quality assurance without continuous developer supervision.
4) Autonomous DevOps and Deployment: Agentic AI manages deployment strategies, monitors system health, automatically scales resources, and executes zero-downtime rollouts, optimizing operations based on usage and performance data. For instance, in fintech applications, agentic AI can detect vulnerable dependencies and patch them in real time before releasing updates.
5) Dynamic Documentation Generation and Maintenance: Maintaining up-to-date technical documentation is a persistent challenge. Agentic AI auto-generates API references, user guides, and system diagrams that evolve alongside code changes, improving maintainability and onboarding.
6) Predictive Performance Optimization: Continuous monitoring by AI agents allows proactive identification and mitigation of performance bottlenecks before they affect end users. This includes adaptive configuration tuning learned from historical system behavior.

©IJRASET: All Rights are Reserved

III. CURRENT STATE OF AI IN PRODUCT SUPPORT

A. *Observability and Predictive Alerting*

Modern IT environments rely heavily on observability platforms to collect comprehensive health and performance metrics from distributed software systems. These platforms ingest telemetry data including logs, traces, and system metrics, enabling continuous monitoring of system behavior. However, traditional threshold-based alerting often leads to noisy alarms and delayed issue detection.

Agentic AI-powered observability enhances this monitoring by learning baseline system patterns and detecting anomalous behavior with greater precision. Machine learning algorithms analyze historical and real-time data to identify subtle deviations indicative of potential failures, such as resource exhaustion or abnormal latency. For example, if disk usage consistently exceeds 80% for a sustained timeframe, the AI agent can generate a prioritized alert for administrators, enabling them to initiate preventive actions like scaling storage capacity or redistributing workloads.

Advancements in AI observability also mitigate alert fatigue by dynamically adjusting alert thresholds and filtering false positives. This refinement allows IT teams to focus only on actionable, high-priority incidents, improving operational efficiency and reducing response times. Furthermore, observability systems equipped with predictive analytics forecast emerging issues before end-user impact, transitioning IT operations from reactive troubleshooting to proactive management.

B. *Toward Autonomous Remediation*

While predictive alerting has improved operational awareness, agentic AI aims to bridge the gap toward autonomous remediation. Future AI agents will not just surface potential problems but also execute predefined or learned corrective actions without human intervention. Leveraging historical incident data and administrator responses, these intelligent systems develop decision models to autonomously mitigate or resolve common failure modes, thereby minimizing system downtime.

For instance, if an AI agent observes consistent patterns where high CPU utilization precedes service degradation, it may automatically invoke resource allocation workflows or recycle failing services. This capability reduces the cognitive load on IT operators and accelerates mean time to resolution (MTTR). Recent developments include AI agents capable of orchestrating complex intervention sequences, interacting with multiple subsystems to restore functionality.

By integrating intelligent automation, AI agents facilitate self-healing systems that adapt in real-time to changing conditions and incidents. This shift enhances system reliability, scalability, and availability, while allowing human operators to focus on strategic improvements and innovation.

IV. FUTURE AGENTIC AI OPPORTUNITIES IT OPERATIONS AND SUPPORT

AI agents play a pivotal role in bridging the gap between customer-site IT administrators and the product vendor’s support engineers, facilitating seamless, collaborative problem resolution. These agents possess advanced contextual understanding capabilities, enabling them to differentiate between transient anomalies and genuine system threats. For instance, a latency spike is not merely treated as momentary noise; rather, the AI agent interprets it as a potential precursor to a cascading system failure, warranting immediate attention. Similarly, an unusual login event is analyzed holistically against known attack patterns, allowing the agent to swiftly identify and escalate possible security breaches.

Beyond recognition, these AI agents proactively initiate remediation measures aimed at forestalling issue escalation. By integrating contextual insights with predictive analytics, they can autonomously execute predefined mitigation protocols or recommend precise corrective actions to administrators. This proactive stance drastically reduces downtime and improves system reliability, ensuring that critical incidents are addressed before impacting end users. By automating these complex diagnostic and intervention workflows, AI agents empower human operators to focus on higher-order tasks, enhancing operational efficiency and strengthening security posture in dynamic IT environments.

This is illustrated in Fig. 1, where multiple agents deployed at customer sites are managed by a central agent hosted at the vendor's support engineer location. The potential developments managed by agents at both locations are detailed here.

A. *AI Assistance for IT Administrators*

1) Root Cause Analysis (RCA) and Automated Decision Capability (ADC): For every system alert, agentic AI can analyze logs, traces, and metrics to offer root cause summaries, enabling administrators to assess the accuracy (true positive/false positive) of alerts and plan corrective action. By monitoring administrator actions, AI agents develop models for automated decision capability, ultimately recommending or performing remediation based on historical and real-time data.

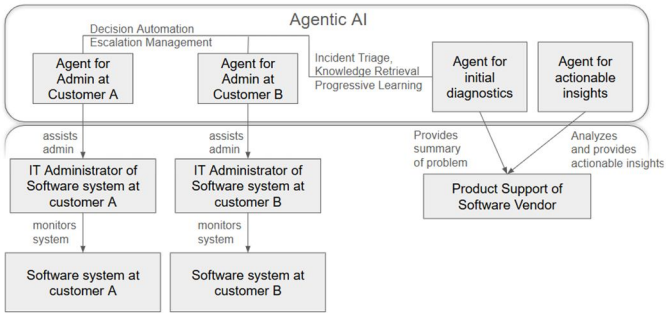


Fig -1: Agentic AI in Product Support

2) Integrated Tool Invocation: Advanced agents can be equipped with a toolkit to perform routine actions—such as log rotation, memory dump collection, or network diagnostics—triggered by detected incidents or RCA outcomes. This capability forms the foundation of an automated operational support cycle.

3) Escalation and Vendor Support Automation: In persistent or complex incidents, AI agents can initiate escalation procedures: filing vendor support tickets, collating diagnostics, and summarizing contextual information for downstream support engineers. By automating these steps, AI relieves administrators of repetitive and time-consuming processes, improving overall incident response efficiency.

B. AI Assistance for Support Engineers

1) Automated Diagnostics Gathering

Upon receiving customer-reported incidents, AI agents immediately parse the initial defect summaries provided by users or monitoring systems, extracting critical keywords, contextual details, affected components, and error signatures. Leveraging natural language processing (NLP) and pattern recognition, these agents often interpret unstructured or semi-structured inputs to construct a comprehensive incident profile. Subsequently, the agents trigger a suite of automated diagnostic tools designed to capture all relevant system metrics, logs, traces, and telemetry necessary for thorough root cause investigation. This may include activating debug logging, memory profiling, network traffic analysis, configuration state captures, or execution of prebuilt diagnostic scripts tailored to the incident type. Such proactive and context-aware diagnostics ensure that support engineers are equipped with a holistic set of high-fidelity data, significantly reducing back-and-forth communications requesting additional information.

By automating these traditionally manual, labor-intensive data collection steps, AI agents accelerate incident resolution workflows, minimize human error, and enable rapid triage. Furthermore, this early comprehensive data capture facilitates sophisticated automated root cause analysis algorithms, which correlate multi-dimensional datasets to pinpoint underlying faults.

Additional benefits of this AI-driven diagnostic orchestration include standardized data collection across incidents, improved reproducibility of investigations, and the ability to continuously refine diagnostic strategies based on incident outcomes and feedback. Ultimately, the integration of intelligent AI agents in diagnostic data capture transforms incident management from reactive troubleshooting into a proactive, accelerated, and highly efficient process.

2) Knowledge Base Search and Retrieval-Augmented Generation (RAG)

Support agents frequently rely on extensive historical records of defect resolutions, troubleshooting guides, and product documentation to efficiently diagnose and resolve customer issues. However, the sheer volume and complexity of such knowledge bases often pose challenges in quickly locating the most relevant information. Recent advances in AI, particularly Retrieval-Augmented Generation (RAG), have revolutionized this process by combining powerful retrieval mechanisms with generative language models to deliver precise, context-aware responses directly to support agents.

RAG models function by first retrieving relevant documents or prior support cases from large organizational knowledge repositories based on the current customer issue query. Unlike standalone language models that depend solely on pre-trained data, RAG dynamically accesses up-to-date and domain-specific information, ensuring accuracy and reducing errors due to outdated or incomplete knowledge. This retrieved information is then incorporated into the generative process, enabling the model to produce detailed, relevant, and coherent explanations or recommendations tailored to the query.

In practical support workflows, RAG facilitates rapid identification of known issues, suggested fixes, and diagnostic procedures by synthesizing past experiences and technical manuals. This dramatically shortens the mean time to resolution by enabling support engineers to leverage cumulative organizational intelligence without manually searching disparate sources. Studies, including implementations at major tech firms, have demonstrated that RAG-based knowledge retrieval can reduce resolution times by approximately 25-30% while increasing first-contact resolution rates and customer satisfaction.

Furthermore, RAG-enhanced systems continually learn from new cases, feedback, and outcomes, refining retrieval relevance and response quality over time. This adaptive capability ensures that support ecosystems remain current with evolving product features and incident patterns.

3) SME Guidance and Progressive Learning

Complex or novel incidents often require the expertise of subject matter experts (SMEs) to diagnose and devise effective resolutions. Agentic AI systems play a crucial role by observing and learning from the investigative strategies employed by engineers and the advice provided by SMEs. Through continuous interaction and feedback, these AI agents gradually build advanced troubleshooting heuristics and decision models tailored to specific domains.

This learning enables AI to handle increasingly complex diagnostic tasks autonomously, reducing the volume of repetitive or routine issues that reach SMEs. By offloading such operational workloads to AI, SMEs can dedicate more time and cognitive resources to high-value innovation, strategic problem-solving, and the development of new methodologies.

Moreover, SMEs actively contribute to the refinement of AI models by validating outputs, correcting errors, and injecting domain-specific nuances that cannot be easily captured by data alone. This synergistic partnership fosters a continuous improvement cycle where AI amplifies SME knowledge while SMEs guide AI evolution.

4) Customer Collaboration and Advanced Diagnostics

AI agents play an instrumental role in enhancing customer interactions during incident management by actively recommending or autonomously executing additional diagnostics tailored to the evolving nature of each incident. These diagnostics may include enabling debug logging, conducting memory profiling, collecting core dumps, or performing targeted data extraction from various subsystems. By dynamically selecting and orchestrating these diagnostic activities, AI agents ensure that relevant, high-quality data is gathered to expedite root cause analysis.

Crucially, these agents do not operate on static scripts but continuously adapt their diagnostic strategies based on incident context changes, historical patterns, and real-time system feedback. For instance, if initial diagnostics indicate intermittent network latency, the AI may initiate deeper packet capture or bandwidth analysis tools. Conversely, if memory leaks are suspected, it might trigger profiling tools optimized for heap and garbage collection metrics. This adaptive intelligence ensures optimal diagnostic coverage while minimizing unnecessary data collection and system overhead.

5) Insight Generation and Action Recommendation

As AI models in IT support and incident management mature, they increasingly prioritize delivering actionable insights and suggesting precise fixes grounded in accumulated diagnostic data and learned patterns of system behavior. By synthesizing rich telemetry, historical resolution records, and contextual information, these advanced AI agents help support engineers focus on the most impactful next steps, enabling rapid and accurate resolution of high-complexity defects that were previously daunting or time-consuming. A critical enhancement in mature systems is the embedding of explainability features, often referred to as Explainable AI (XAI). XAI provides transparency into how and why AI agents generate specific recommendations or predictions, breaking down decision factors in a human-understandable manner. This empowers support engineers and administrators to validate AI suggestions confidently and improves trust in automated insights. Moreover, this transparency helps teams identify false positives or biases, refining AI accuracy and reliability continuously.

Human-in-the-loop (HITL) controls complement this approach by allowing experts to review, override, or fine-tune AI-driven actions before execution, striking a balance between automation efficiency and human judgment. This critical oversight safeguards against potential errors, reduces operational risks, and ensures compliance with organizational security policies and regulations.

Incorporating explainable, auditable AI systems also addresses increasing regulatory demands for accountability and traceability in automated decision-making processes, especially in sensitive IT infrastructure contexts. These capabilities contribute to more robust cybersecurity postures, as AI agents transparently reveal suspicious patterns or anomaly triggers that inform verification and response strategies.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Ultimately, the combination of mature AI-driven prioritization, transparent decision-making, and human-guided controls empowers support teams to accelerate problem resolution, mitigate risks proactively, and maintain operational resilience in increasingly complex technological landscapes.

V. TAXONOMY OF SUPPORT-CENTRIC AI AGENT TASKS

We propose the following categorization for support-oriented agentic AI:

1) Decision Automation: Suggesting or executing fixes
2) Escalation Management: Vendor communication and ticketing
3) Incident Triage: Automated summary and initial data gathering
4) Knowledge Retrieval: Rapid reference to prior solutions
5) Progressive Learning: Model evolution from expert interactions

VI. NAVIGATING THE ADOPTION JOURNEY: A STRUCTURED 5-STEP FRAMEWORK

Adopting an agentic system of this nature requires deliberate planning and execution, guided by a structured 5-step framework to ensure successful integration and operational effectiveness.

Failure to address these areas may result in AI systems that optimize for efficiency at the cost of fairness, empathy, or financial controls. Consequently, Chief Information Officers(CIO) and IT leaders must carefully balance autonomy with rigorous accountability.

A. Identify High-Impact Use Cases

Begin by targeting areas characterized by repetitive, high-volume tasks that yield measurable operational benefits. Examples include log triage, password resets, and configuration rollbacks. Focusing on these high-impact processes ensures early value realization and resource optimization.

B. Start with Value-First Pilots

Deploy small-scale pilot projects to validate the return on investment (ROI) before committing to larger roll out. Select scenarios that offer quick, demonstrable wins to build organizational confidence and encourage stakeholder buy-in through tangible results.

C. Establish Governance and Guardrails

Develop and enforce policies governing data access, usage transparency, and ethical AI practices. Incorporate human oversight mechanisms where necessary to maintain accountability and ensure compliance with internal standards and external regulations, fostering trustworthy AI deployment.

D. Scale with Continuous Learning

Expand the agent's autonomy incrementally as organizational confidence and system maturity increase. Integrate feedback loops that incorporate new data and stakeholder insights to continuously enhance and refine model performance and reliability.

E. Addressing the Pitfalls: Data, Security, and Culture

1) Data Quality: High-quality, well-governed data pipelines are foundational; poor data leads to suboptimal decision-making and risks system integrity.
2) Security & Compliance: AI models' opaque decision-making processes pose regulatory compliance challenges. Ensuring explainability and establishing proactive monitoring frameworks are critical to managing risks and meeting audit requirements.
3) Culture of Innovation: Resistance from teams wary of ceding decision power to AI systems is common. Cultivating transparency, up-skill personnel, and implementing AI literacy initiatives are essential steps to build trust and drive adoption.

VII. BUSINESS IMPACT: MEASURING ROI OF AGENTIC AI

Agentic AI extends beyond automating routine operations; it delivers following quantifiable business outcomes to measure the Return On Investment (ROI):

1) Cost Savings: Automated incident triage and remediation reduce operational overhead and free up skilled personnel for higher-value tasks.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

2) Reduced Mean Time to Repair (MTTR): Accelerated incident resolution enhances system uptime and bolsters customer satisfaction by minimizing service disruptions.
3) Improved Service Quality: Predictive insights enable proactive disruption avoidance, leading to a more seamless and reliable user experience.
4) Strategic Agility: By offloading repetitive workload, IT teams can redirect their focus towards innovation, digital transformation, and strategic initiatives.

Organizations adopting Agentic AI today position themselves as resilient, cost-efficient, and innovation-ready leaders in tomorrow's competitive landscape.

VIII. CHALLENGES AND FUTURE DIRECTIONS

Despite the significant promise demonstrated by agentic AI systems in automating IT support and operations, several critical research and engineering challenges must be addressed to realize their full potential:

1) Data Quality, Accessibility, and Labeling: Building robust operational AI models depends heavily on access to high-quality, well-labeled data. Many organizations face challenges obtaining comprehensive datasets due to fragmented systems, inconsistent logging practices, and privacy constraints. Developing scalable methods for data curation, anonymization, and labeling remains a high priority.
2) Balancing Automation with Human Oversight and Regulatory Compliance: While automation improves efficiency, maintaining appropriate levels of administrator control is essential to prevent unintended consequences, ensure accountability, and comply with increasingly stringent regulations around data privacy and operational safety. Designing flexible systems that empower human-in-the-loop intervention while maximizing autonomous capabilities is a complex challenge.
3) Ensuring Explainability and Transparency: As AI-driven decisions impact mission-critical environments, explainability becomes vital for trust, auditing, and debugging. Current AI models, particularly deep learning ones, often operate as "black boxes," hindering acceptance by human stakeholders. Advancing methods that offer interpretable insights into AI reasoning without compromising performance is necessary.
4) Advancing Secure Tool Automation: Deploying AI agents that directly execute actions in production systems introduces substantial security risks, including potential exploits and unauthorized escalations. Research into secure, granular access controls, verifiable execution, and fail-safe mechanisms is critical to safeguarding sensitive environments.
5) Integration and Interoperability: Agentic AI solutions must seamlessly integrate with legacy infrastructures and diverse enterprise ecosystems, which is often hindered by incompatible APIs, data silos, and operational fragmentation. Architecting modular, vendor-agnostic frameworks with standardized interfaces will facilitate broader adoption.
6) Managing Systemic Risks and Agent Sprawl: The increasing number and autonomy of AI agents can lead to operational chaos without centralized orchestration, visibility, and control. Best practices for agent governance, life cycle management, and minimizing redundant or conflicting autonomous actions are emerging research areas.
7) Ethical and Bias Considerations: Mitigating bias in AI-driven operational decisions and ensuring fairness remain ongoing concerns. Proactive bias detection, ethical guidelines, and continuous monitoring are essential for responsible AI deployment.

Addressing these challenges requires interdisciplinary collaboration among academia, industry, and regulatory bodies. Future research should focus on developing trustable, transparent, and secure agentic AI architectures capable of autonomous yet controllable decision-making in complex IT environments. Through deliberate design and rigorous evaluation, agentic AI can transition from promising prototypes to reliable, scalable tools that revolutionize IT product support.

IX. CONCLUSION

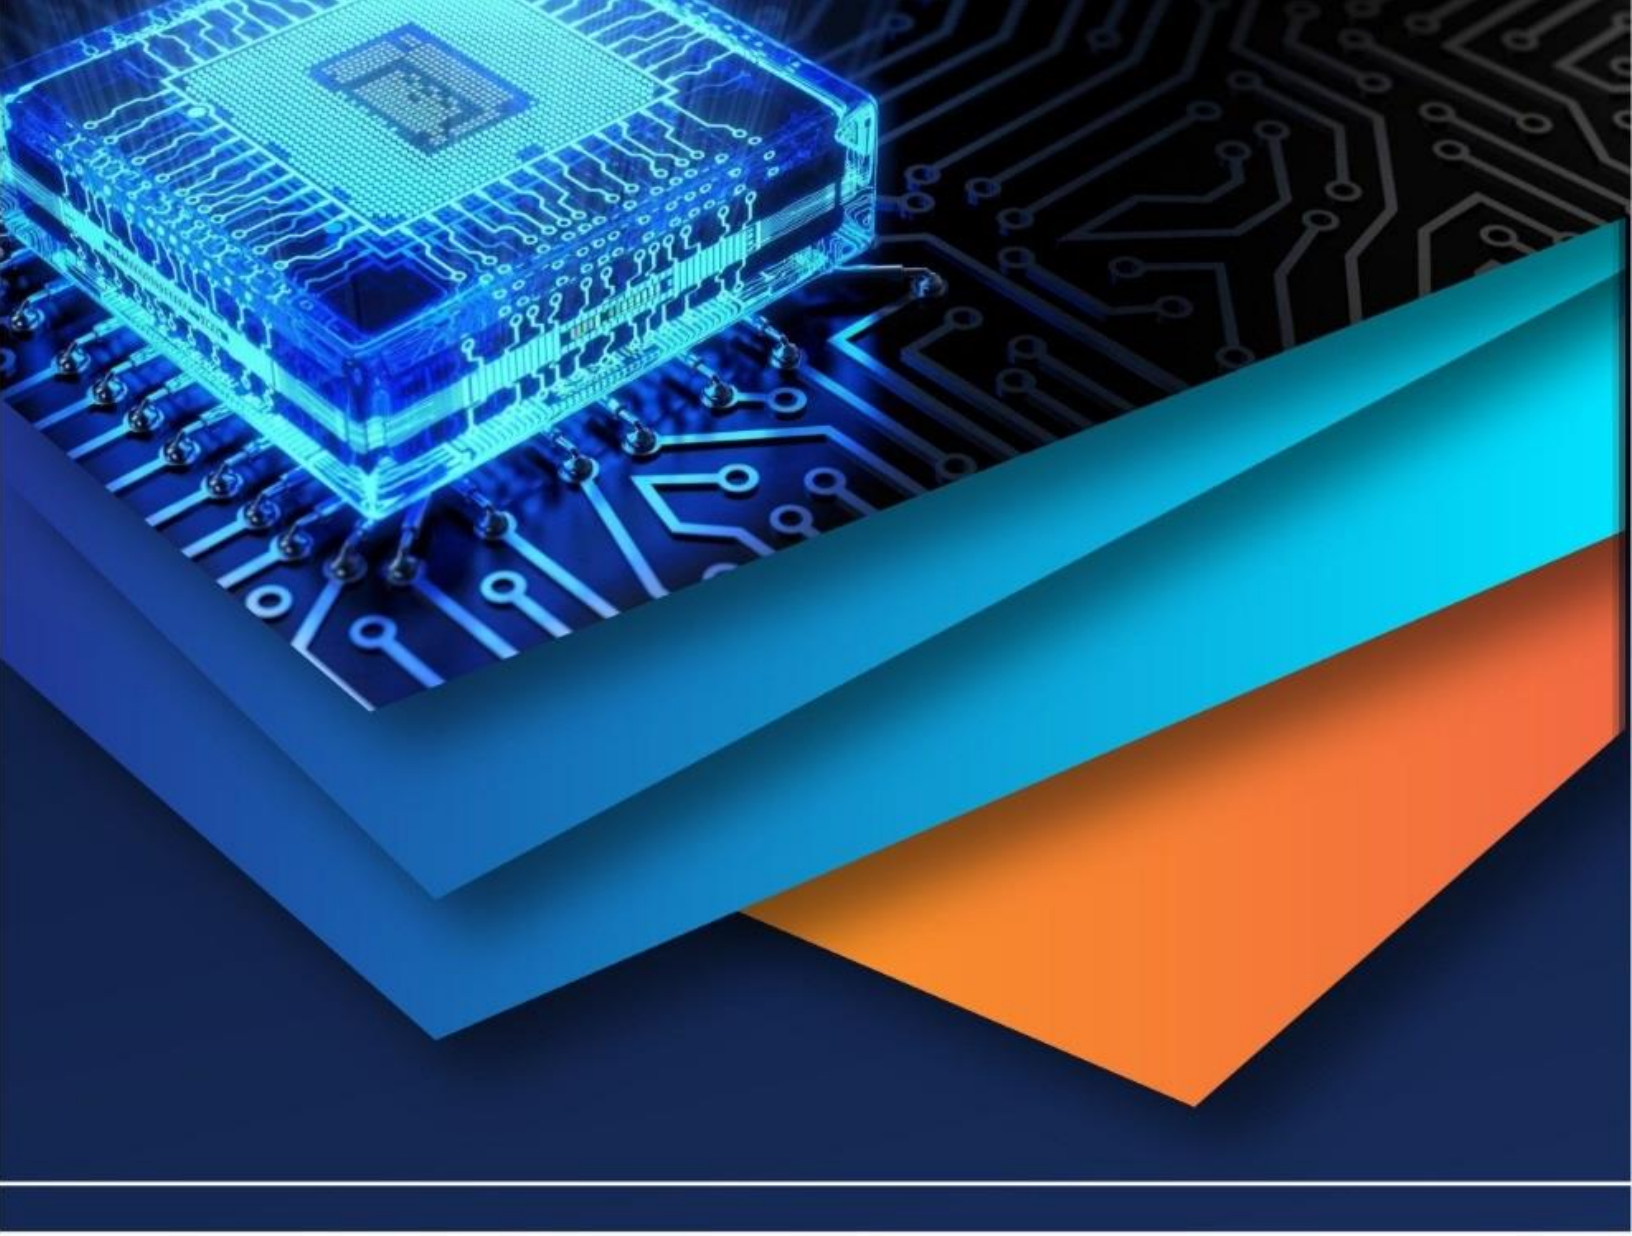
Agentic AI systems are redefining the future of software development and operational support by automating routine activities, expediting diagnosis, and augmenting decision-making. The road map outlined in this paper underscores the trans formative potential—and ongoing challenges—of deploying AI agents in support engineering. By evolving from coding assistants to autonomous operational partners, agentic AI will ultimately drive efficiency, reliability, and resilience across the software product life cycle.

REFERENCES

[1] N. Kshetri, "Revolutionizing Higher Education: The Impact of Artificial Intelligence Agents and Agentic Artificial Intelligence on Teaching and Operations," in IT Professional, vol. 27, no. 2, pp. 12-16.

[2] V. Tupe and S. Thube, "Demonstrating Multi-Agent Collaboration via Agent-to-Agent and Model Context Protocols: An IT Incident Response Case Study," 2025 IEEE International Conference on Service-Oriented System Engineering (SOSE), Tucson, AZ, USA, 2025, pp. 1-5
[3] Tanveer Aamina, Mohammed Zaid, Syeda Huda."Evaluating Multi-Agent AI Systems for Automated Bug Detection and Code Refactoring", Volume 13, Issue X, International Journal for Research in Applied Science and Engineering Technology (IJRASET) Page No: 12-21, ISSN : 2321-9653, www.ijraset.com.
[4] Tanveer Aamina, Mohammed Zaid, Syeda Huda."Evaluating Multi-Agent AI Systems for Automated Bug Detection and Code Refactoring", Volume 13, Issue X, International Journal for Research in Applied Science and Engineering Technology (IJRASET) Page No: 12-21, ISSN : 2321-9653, www.ijraset.com
[5] Tanveer Aamina, Mohammed Zaid, Syeda Huda."Evaluating Multi-Agent AI Systems for Automated Bug Detection and Code Refactoring", Volume 13, Issue X, International Journal for Research in Applied Science and Engineering Technology (IJRASET) Page No: 12-21, ISSN : 2321-9653, www.ijraset.com

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)