Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Autonomous Incident Remediation via GenAI-Assisted Runbooks

Saravanan Raj

Independent Researcher, USA

ARTICLE INFO

Received: 30 Sept 2025

Revised: 05 Nov 2025

Accepted: 13 Nov 2025

ABSTRACT

Cloud infrastructure has evolved to form the backbone of global digital operations; however, reliability engineering has not kept pace with this growing complexity. While observability and alerting systems have matured considerably, incident response still relies heavily on human expertise, thus introducing a significant gap between detection and automated remediation that translates into very costly downtime and operational fatigue. This article introduces a closed-loop remediation framework powered by Generative AI, enabling the diagnosis, execution, and validation of incident resolution with appropriate safety guardrails and auditability. Within this framework, the model integrates LLM-based diagnostics, policy-driven execution, and safety validation mechanisms, with continuous learning feedback cycles. The article draws best practices from multi-cloud implementation and puts forward a phased implementation approach while discussing governance considerations toward autonomous remediation. This framework demonstrates substantial improvement in incident resolution speed and reduction in manual escalations, thus positioning autonomous incident remediation as the cornerstone for the next evolution in AI-powered reliability engineering. As systems continue to scale beyond human cognitive limits, such autonomous approaches are not just advantageous but also become very key for operational resilience amidst ever-increasing digital environment complexities.

Keywords: Generative AI, Autonomous Remediation, Cloud Reliability, Closed-Loop Framework, Policy-Driven Automation

1. The Evolution of Cloud Reliability Through Generative AI

Cloud reliability engineering has become a cornerstone for the digital economy, with financial systems, healthcare networks, and more that rely on resilient cloud environments. For this reason, both cloud infrastructure and AI operation have been listed among several critical technologies that are considered to be significant to national security, by the White House Office of Science and Technology Policy [1].

Although the SRE has been improved, and tools that assume observability have been developed, incident management remains an important operational issue. Organizations indicate that detection and remediation processes remain limited by disjointed technological ecosystems and manual escalation procedures. Such incompatibility between surveillance opportunities and resolving efficiency leaves a gap in operations. The microservice architecture, deploying to multiple clouds, and edge computing complicates the provision of responses, and human-centered response models are no longer sufficient to guarantee service levels.

Traditional incident management approaches have particular limitations in high-change-velocity environments. The thousands of code changes a week introduced via modern deployment pipelines create a dynamic landscape that no longer allows for static runbook approaches. Traditional automation systems are based on deterministic pattern matching that does not work well with novel failure scenarios; hence, diverting engineering resources to reactive troubleshooting rather than proactive improvements.

Generative Artificial Intelligence redefines incident management with its unique combination of contextual analysis and natural language processing. Unlike other rule-based systems, GenAI processes heterogeneous data formats in parallel, extracting meaning from logs, metrics, distributed traces, and configuration repositories. That would also facilitate the establishment of correlations

between seemingly unrelated signals that human analysis might miss. The technical architecture is usually based on retrieval-augmented generation, combining real-time observability data with historical incident knowledge to ground the analysis in the contexts of an organization while using general principles of troubleshooting.

Autonomous remediation architecture must be developed with great care so that it can maintain governance and safety. More sophisticated versions apply multistage pipelines, separating the diagnostic analytics from the execution. The diagnostic element formulates causal hypotheses from the telemetry data, while the execution element applies a policy matrix according to the criticality of service and the scope of impact. This separation ensures that automated actions remain bound by organizational policy rather than emergent model behaviors.

GenAI-powered reliability systems build feedback loops into the core of continuous learning and improvement. In the case of successful remediation, this cements in place associations between symptoms and interventions; when an action does not succeed, the system revises its causal models. This makes for a virtuous circle where operational knowledge accrues in a common system rather than solely within a responder's experience.

Other challenges in implementation include data quality, as these systems need large historical incident datasets with well-documented symptoms, actions, and outcomes. Integration with infrastructure-as-code systems requires API design with due care to keep the declared state of the system consistent.

With cloud environments now more extensive in scale and criticality, the autonomous reliability systems represent not just efficiency gain but a strategic necessity. Further evolution points toward more sophisticated AI-human collaboration, where automated systems handle routine incidents, while human expertise addresses novel failure modes and architectural improvements. Organizations that successfully deploy advanced architectures of reliability will create sustainable advantages via superior service continuity and more efficient use of engineering resources [2].

2. The Automation Gap in Current Incident Management

The current incident management systems have highly automated portions of the response procedure; nonetheless, automation-complete and accurate diagnosis, correct mitigation, and extensive verification, including human intervention-free, is elusive in most business contexts. This is a gap that is one of the largest challenges that organizations encounter in the optimization of operational resilience and maximum engineering resource optimization.

The fragmentation of operational toolchains is a chief barrier to fully autonomous remediation. Most enterprise architectures have distinct systems for monitoring, alerting, ticketing, and the execution of automation, creating artificial boundaries around related data for operations. Segmentation prevents the kind of holistic context required for confident automated decision-making. Without telemetry data residing in the same place as historical incident records, and the execution platforms of runbooks integrated with configuration management systems, the context is fragmented, and synthesis requires crossing system boundaries performed by humans. Forsgren and colleagues have shown that tool fragmentation strongly leads to longer incident resolution times across a wide range of industries [3].

Governance gaps further hinder the adoption of autonomous remediation. Organizations often lack standardized risk classification frameworks that inform which automated actions can execute automatically versus those that require human approval. Without structured governance around what remediations can safely run automatically versus which ones require human approval, organizations default to conservative approaches wherein most scenarios require manual intervention. This governance gap plays out in ad-hoc approval workflows that vary by team, technology stack, and individual service, in a way that creates an environment where automation boundaries appear

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

arbitrary rather than risk-calibrated. The lack of formalized risk classification frameworks has been called out as one of the primary barriers to autonomous operations adoption in multiple industry surveys.

The auditability gap poses another major barrier to autonomous remediation. Incident response systems in the present day may not fully track the automated action process, so it is hard to follow the chain of decisions that culminates in a specific intervention. This shortcoming causes high compliance risks and particularly in controlled sectors where the control of operations must be provable. Equally, autonomous systems cannot be used to satisfy regulatory requirements without transparent, unalterable audit trails of not only the rationale behind an automated decision, but also the actual actions taken. This compliance risk leads organizations to maintain human approval gates where proper audit mechanisms would allow for their elimination, as noted in recent literature on technical governance [4].

The cognitive burden on on-call engineering personnel represents perhaps the most immediate consequence of insufficient automation. Typical enterprise alert volumes have grown exponentially with the proliferation of microservices and distributed architectures, routinely overwhelming human processing capacity. Engineering teams face alert fatigue from hundreds of daily notifications, many representing duplicative symptoms of the same underlying issue. This cognitive overload leads to increased error rates, delayed responses to critical alerts, and substantial quality-of-life impacts for technical personnel. The psychological impact of excessive on-call burden has been linked to engineering burnout and turnover, creating significant organizational costs beyond the immediate incident impact.

Despite considerable investments in automation technologies, incident response workflows continue to involve humans at most steps of the process. This pervasive inefficiency is directly related to the economic bottom line because of longer outage durations, engineering resources, and opportunity costs, since technical teams are focused on operational firefighting rather than innovation. For example, in organizations where the incident response tasks could be theoretically automated, a large fraction of the capacity of engineering teams reportedly goes into their manual execution, acting as a drag on development velocity and technical improvement.

The incident management automation disparity is not merely a technical issue but a strategic constraint on the efficiency and resilience of a company. To solve this gap, it is necessary, in addition to the technological solutions, to have governance structures, audit procedures, and cultural modulations that will open the way to genuinely autonomous operations. As the complexity of the system increases, therefore, bridging this gap in automation becomes essential to sustainable operational models.

Barrier	Impact	Potential Solution
Toolchain Fragmentation	Extended incident resolution times	Unified observability platforms
Governance Deficiencies	Conservative automation approaches	Standardized risk classification frameworks
Auditability Gaps	Compliance challenges in regulated sectors	Transparent, immutable audit trails
Cognitive Burden	Alert fatigue and engineering burnout	Intelligent alert correlation and prioritization
Cultural Resistance	Maintenance of human approval gates	Phased automation adoption with proven outcomes

Table 1: Impact of Automation Gaps on Operational Efficiency [3, 4]

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

3. Generative AI's Role in Operations

Generative Artificial Intelligence, with Large Language Models (LLMs) as its primary implementation, has fundamentally transformed operational capabilities in cloud environments by enabling contextual reasoning rather than merely pattern classification. Traditional operational tools relied on predefined rules and signature-based detection, creating brittle systems that could not adapt to novel failure modes. In contrast, modern GenAI systems demonstrate emergent capabilities for synthesizing disparate operational signals into coherent diagnostic narratives, representing a paradigm shift in incident management approaches.

The telemetry interpretation capabilities of GenAI systems provide unprecedented operational context awareness. These models can analyze complex time-series anomalies across distributed systems, correlating current observations with historical incident patterns to identify root causes with minimal human guidance. This capability extends beyond simple threshold violations to recognize subtle pattern shifts that might indicate impending system degradation. By maintaining a semantic understanding of system architecture and component relationships, GenAI systems can distinguish between coincidental correlations and causal factors—a distinction that traditional monitoring systems struggle to establish. Research by Schmidt and colleagues demonstrates that GenAI-based anomaly detection achieves significantly higher precision and recall metrics compared to traditional statistical approaches across diverse operational datasets [5].

Natural language interfaces represent another transformative operational capability enabled by GenAI. These systems can generate contextually appropriate remediation commands based on diagnostic findings, presenting them in human-readable formats that on-call engineers can review before execution. This capability bridges the semantic gap between observation and action, transforming raw telemetry into executable intervention strategies. For organizations with complex technology stacks requiring specialized domain knowledge, these generated remediation suggestions democratize operational expertise by encoding best practices into accessible language. The bidirectional translation between machine telemetry and human-readable instructions significantly reduces the expertise barrier for effective incident response.

Status summarization capabilities address the cognitive overload challenges inherent in complex incident response. GenAI systems can distill hundreds of log entries, metrics, and alerts into concise situation assessments that highlight critical factors while filtering noise. This capability proves particularly valuable for incidents requiring escalation approval, as it enables rapid knowledge transfer between technical teams and decision-makers. By producing multi-level summaries tailored to different stakeholder perspectives, these systems facilitate informed decision-making without overwhelming participants with excessive technical detail.

The continuous learning capacity of GenAI systems represents perhaps their most significant operational advantage. Through systematic feedback integration, these systems refine their diagnostic accuracy and remediation recommendations based on observed outcomes. Unlike traditional static runbooks, GenAI systems continuously incorporate new failure modes and successful remediation strategies, creating an adaptive knowledge base that evolves with the underlying technology landscape. This self-improving capability allows organizations to capture operational knowledge that traditionally remained siloed within individual experts, transforming it into an organizational asset that benefits all responders.

Recent technical advances in LLM-based diagnostics have enhanced these capabilities through retrieval-augmented generation (RAG) architectures. By combining the semantic reasoning capabilities of language models with structured knowledge retrieval from operational databases, RAG systems achieve precision that was previously unattainable. These systems can seamlessly integrate structured telemetry data (metrics, service health indicators) with unstructured information (logs, incident notes, documentation) into a unified analytical framework. The resulting insights benefit

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

from both the pattern recognition capabilities of traditional monitoring and the contextual understanding of language models.

The transition from conventional AIOps implementations to fully autonomous remediation systems necessitates rigorous governance frameworks built on three foundational principles. Policy control mechanisms must establish explicit boundaries for automated actions, classifying operations by risk level and defining appropriate approval workflows for each category. These policies must reflect organizational risk tolerance and compliance requirements, ensuring that automation operates within defined safety parameters. Technical implementations of these policy frameworks typically involve rule engines that evaluate proposed actions against predetermined criteria before execution authorization.

Safety validation and rollback capabilities constitute the second essential principle for autonomous remediation. Every automated intervention must include pre-execution validation checks that verify system readiness, execution monitoring that detects unexpected outcomes, and rollback mechanisms that can restore previous states if necessary. These safety systems must operate at the same automation level as the primary remediation actions, ensuring that recovery can proceed even if the original automation fails. The implementation of comprehensive validation frameworks has been identified as a critical success factor in autonomous operations adoption, as detailed in technical surveys of enterprise automation practices [6].

Comprehensive auditability represents the third non-negotiable principle for autonomous remediation systems. Every automated action must generate immutable audit records documenting the observed symptoms, diagnostic reasoning, decision criteria, execution steps, and validation results. These audit trails must be machine-readable for automated analysis while remaining interpretable by human reviewers for governance purposes. The granularity of these records should enable precise reconstruction of the decision sequence leading to any automated intervention, satisfying both operational troubleshooting needs and compliance requirements.

As cloud operations continue to increase in complexity and scale, the role of GenAI in enabling autonomous remediation will only expand. Organizations that successfully implement these technologies with appropriate governance frameworks will achieve significant competitive advantages through enhanced reliability, reduced operational overhead, and more efficient utilization of engineering expertise.

Capability	Traditional Systems	GenAI Systems	Key Benefit
Diagnostic Intelligence	Pattern/rule-based	Contextual reasoning	Adapts to novel failure modes
Telemetry Analysis	Threshold violations	Semantic understanding	Higher precision/recall in anomaly detection [5]
Interface Mechanisms	Raw data outputs	Natural language commands	Democratized operational expertise
Information Processing	Manual correlation	Automated summarization	Reduced cognitive overload
Knowledge Management	Static runbooks	Continuous learning	Evolving, organization-wide expertise
Integration Approach	Siloed tools	RAG architecture	Unified analytical framework

Table 2: Comparative Advantages of GenAI vs Traditional Monitoring [5, 6]

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

4. The Closed-Loop Remediation Framework

The safe operationalization of Generative Artificial Intelligence (GenAI) in cloud reliability engineering requires structured governance mechanisms that balance autonomy with control. The Closed-Loop Remediation Framework provides a comprehensive architectural model for implementing autonomous incident resolution while maintaining appropriate safety boundaries and audit capabilities. This four-stage framework integrates AI-powered diagnostic capabilities with policy governance, validation controls, and continuous learning mechanisms to create a self-improving system for incident remediation.

4.1 Stage 1: AI-Based Diagnostics

The diagnostic foundation of the framework leverages Large Language Models (LLMs), specifically fine-tuned on historical incident data, to perform contextual analysis of operational telemetry. Unlike traditional rule-based diagnostics, these models can synthesize heterogeneous data types—structured metrics, semi-structured logs, and unstructured trace information—into coherent causal hypotheses. The diagnostic process begins with symptom characterization, then progresses through correlation analysis to identify potential root causes and their probability distributions. The model's ability to interpret semantic relationships between disparate signals allows it to recognize subtle patterns that would elude conventional monitoring systems.

The technical implementation typically involves a retrieval-augmented generation (RAG) architecture where the language model interfaces with knowledge bases containing service topologies, historical incidents, and documented failure modes. The diagnostic output includes both root cause hypotheses and associated confidence metrics, ensuring transparency in the reasoning process. For each identified cause, the system generates candidate remediation steps based on historical efficacy data and current system state, prioritizing interventions with the highest probability of resolving the incident with minimal collateral impact. Research by Khattak and colleagues demonstrates that fine-tuned LLMs achieve significantly higher diagnostic accuracy compared to traditional pattern-matching approaches, particularly for complex, multi-component failures where causal chains span multiple services [7].

4.2 Stage 2: Policy-Based Execution

Once diagnostic analysis produces remediation candidates, the policy-based execution stage applies governance controls to determine appropriate automation boundaries. This stage implements a structured risk matrix that classifies potential actions according to their impact scope, reversibility, and historical reliability. The policy engine evaluates each proposed remediation against these classification criteria before determining execution permissions.

For services classified as low-risk within the organizational taxonomy, the framework permits full automation of standardized remediation procedures such as service restarts, cache purging operations, and non-disruptive configuration adjustments. These actions typically affect isolated components with minimal dependency chains and established recovery patterns, making them suitable candidates for autonomous execution. Medium-risk services introduce an additional safety layer by requiring rollback preparation before execution. The system automatically creates system state snapshots and validates rollback procedures before implementing remediation actions, ensuring that any unexpected consequences can be rapidly reversed. High-risk services or critical infrastructure components require explicit human approval through defined escalation pathways. The system presents diagnostic findings, proposed remediation steps, and expected outcomes to authorized approvers, who can then authorize execution while maintaining full visibility into the automated process.

This tiered approach to policy governance ensures that automation boundaries reflect organizational risk tolerance and compliance requirements rather than technological limitations. By embedding risk

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

classification into the execution pathway, the framework provides consistent governance across different services and teams while maintaining appropriate human oversight for critical systems.

4.3 Stage 3: Safety Validation and Audit Control

The safety validation stage implements comprehensive verification mechanisms surrounding every automated remediation action. This multi-phase validation approach begins with pre-execution checks that validate system readiness and dependency status before implementing changes. These checks confirm that target systems are in an expected state, that required dependencies are available, and that any preconditions for successful remediation are satisfied.

During execution, the system maintains real-time monitoring of system health indicators to detect unexpected behavior that might warrant intervention. This continuous validation allows for immediate termination of problematic remediations before they cause cascading failures. Post-execution validation implements a comprehensive set of service checks to verify successful resolution, including availability probes, performance metric validation, and functional testing appropriate to the affected service. If post-checks detect health regression exceeding defined thresholds—such as latency increases, error rate spikes, or availability degradation—the system automatically triggers rollback procedures to restore the previous stable state.

The audit control component maintains comprehensive, immutable records of the entire remediation process. These records document the initial system state, diagnostic reasoning, applied policies, execution details, validation results, and outcome. The granularity of these audit trails allows for precise reconstruction of the decision sequence leading to any automated intervention, satisfying operational troubleshooting needs and compliance requirements. Industry analysis by Gartner indicates that robust audit mechanisms represent a critical success factor in autonomous operations adoption, particularly in regulated industries where governance requirements mandate comprehensive traceability [8].

4.4 Stage 4: Feedback and Continuous Learning

The self-improving nature of the framework emanates from the feedback and continuous learning stage, which forms a virtuous cycle of operational improvement. Each remediation incident, successful or failed, provides the system with structured outcome data to feed back into its training corpus. Successful remediations reinforce the relationships between observed symptoms and effective interventions to thereby improve the diagnostic accuracy of the system in the future. Equally valuable learning opportunities arise from failed remediations that identify ineffective approaches, thus triggering policy refinements.

The mechanisms of the continuous learning process have multiple feedback mechanisms. Supervised fine-tuning adapts the diagnostic model parameters by using the outcome data, successively enhancing the model to detect root causes and recommend effective remedial actions correctly. Policy evolution comes about with the analysis of remediation success rates across different service categories, enabling data-driven adjustments in risk classification and boundaries of automation. This improvement over time reduces false positives as the system identifies and eliminates diagnostic patterns leading to unnecessary remediation attempts.

It is this architecture of learning that creates the foundation of a cumulative learning system based on every incident, and thus transforms experience into an institutional capacity and not an expertise in silos. The result is a remediation framework that demonstrates increasing performance on essential operational metrics such as time to fix, first attempt success rates, and automation coverage.

Closed-Loop Remediation Framework is a major advancement in the field of operational resilience that integrates the high-quality analytical abilities of generative AI with well-organized governance, overall safety mechanisms, and ongoing improvement measures. In addition, organizations will

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

realize the benefits of autonomous remediation while retaining adequate control, visibility, and conformance with the requirements of operational governance.

Stage	Key Function	Technical Implementation	Business Value
1. AI-Based Diagnostics	Symptom characterization to root cause identification	Fine-tuned LLMs with RAG architecture	Higher diagnostic accuracy for complex failures [7]
2. Policy-Based Execution	Risk-based automation governance	Structured risk matrices with tiered permissions	Balanced automation with appropriate controls
3. Safety Validation	Pre/during/post-execution verification	Health monitoring and rollback procedures	Reduced risk of cascading failures [8]
4. Feedback and Learning	Performance improvement cycle	Supervised fine-tuning and policy evolution	Institutional knowledge accumulation

Table 3: Four-Stage Closed-Loop Remediation Process [7, 8]

5. Implementation Best Practices and Future of Autonomous Reliability

5.1 Implementation Best Practices

Successful autonomous remediation requires both technical excellence and strategies for organizational adoption. At its base lies quality runbook engineering, which treats remediation logic as first-class software artifacts rather than as byproducts of operational work. Declarative, infrastructure-as-code approaches have demonstrated better maintainability by expressing the desired state rather than the steps to achieve it, allowing for version control and modular composition. DevOps Research and Assessment finds that organizations that have rigorous automation testing practices are significantly more likely to achieve higher automation success rates and faster incident resolution times.

A risk-graduated approach to implementation helps engender an organizational level of trust incrementally, starting with low-risk domains before going on to critical systems. It allows teams to iteratively refine capabilities while demonstrating tangible improvements in reliability that help overcome operational skepticism through results.

Comprehensive governance frameworks turn the formerly autonomous remediation of perceived risk into organizational assets with clearly defined policy matrices that constrain automation boundaries. Trust can be reinforced by an essential record of the decisions made, actions undertaken, and the results achieved in the validation in the process of auditability.

The models of human-AI collaboration present hybrids that strike a balance between automation efficiency and human supervision. Progressive automation may start with "human-approved" modes, which would require explicit approval before execution, expanding autonomy gradually as confidence develops from successful interventions.

Feedback-driven continuous improvement mechanisms enable systems to learn from each incident. Systematic outcome logging documents resolution success and unexpected consequences, creating a training corpus for model refinement. Studies by the leading IT service management show that the organizations implementing structured feedback mechanisms achieve compounding improvements in key reliability metrics over time [10].

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Implementation is done based on a phased roadmap that consolidates observability data, establishes operational foundations, implements AI diagnostics, integrates policy-driven execution, and embeds feedback mechanisms.

5.2 Future of Autonomous Reliability

Cloud reliability engineering is at an inflection point, with a widening gap between human cognitive capacity and system complexity. The evolution of reactive to autonomous operations is an operational imperative in maintaining resilience. Research from IBM's Autonomic Computing Laboratory shows that human operators can monitor about 200 metrics effectively, while today's distributed applications generate millions of telemetry signals from thousands of components, as seen in [11].

The GenAI-assisted framework provides a pragmatic blueprint, balancing automation with governance controls. By embedding policy boundaries and audit capabilities, organizations achieve AI-driven remediation benefits while retaining the oversight levels required. Feedback learning mechanisms create continuously improving capabilities rather than static implementations. Research by McKinsey on AI-enabled operations suggests that systems with continuous feedback loops show performance improvements that compound over time, with the most mature implementations realizing order-of-magnitude gains in key reliability metrics [12].

This shifts the role of the reliability engineers from reactive troubleshooters to proactive system designers, freeing cognitive burden and allowing them to concentrate their work on architectural improvements. Convergence of traditional reliability engineering with AI creates capabilities for continuous evolution of systems based on operational experience. Hence, critical infrastructure is getting increasingly reliant on intelligent automation; hence, companies have to put efficiency and transparency into a balance. The ones that will successfully cross this transition will mark the next generation of digital services, which will be characterized by a long-lasting, reliable under even the most continuously increasing complexity.

Implementation Factor	Current State	Future Direction	Strategic Impact
Runbook Engineering	Manual scripts	Declarative IaC approaches	Higher success rates [9]
Implementation Strategy	Siloed automation	Risk-graduated adoption	Incremental trust building
Governance Framework	Ad-hoc policies	Comprehensive matrices	Risk-appropriate automation
Human-AI Collaboration	Manual with tool assistance	Progressive autonomy	Balanced control and efficiency
Learning Mechanisms	Static procedures	Feedback-driven improvement	Compounding metric gains [10]

Table 4: Key Factors for Successful Autonomous Remediation Implementation [9, 10]

Conclusion

The autonomous reliability evolution is a paradigm shift in dealing with operational resilience within an organization. With distributed systems increasingly expanding in size and complexity to levels that are beyond the human mind, the proposed remediation framework with GenAI assistance provides a sensible model to follow that does not compromise automation ability but offers the necessary

governance controls. By integrating policy delineations, verification systems, and extensive audit trails into the autonomous systems, in such a way, organizations can achieve the advantages of AI-driven incident management without losing the right safety guardrails. This transformational effect is not limited to operational measures but is destined to change the nature of the reliability engineer from being a reactive troubleshooter to an active system engineer and will offload cognitive load and enhance career satisfaction in equal measures. The combination of conventional reliability engineering and artificial intelligence introduces emergent capability that cannot be accomplished separately, i.e., systems responding to the current incidents and developing over time through systematic learning processes. Since critical infrastructure is becoming more and more dependent on intelligent automation, the operational governance systems that envelop such applications will influence the manner in which society can trust autonomous systems. Organizations that have managed to navigate this transition will shape the next generation of digital services, in which there is sustained reliability in the face of increasingly more complexity, and in which autonomous remediation is a strategic requirement to long-term digital resilience, and not an operational improvement.

References

[1] White House Office of Science and Technology Policy, "Critical and Emerging Technologies List Update," 2024. https://www.govinfo.gov/content/pkg/CMR-PREX23-00185928/pdf/CMR-PREX23-00185928.pdf

[2] Zhuangbin Chen et al., "AIOps Innovations of Incident Management for Cloud Services," 2020. https://cloudintelligenceworkshop.org/2020/content/AIOps%20Innovations%20of%20Incident%20Management%20for%20Cloud%20Services.pdf

[3] Nicole Forsgren, Jez Humble, and Gene Kim, "Accelerate: The Science of Lean Software and DevOps Building and Scaling High-Performing Technology Organizations," ACM Digital Library, 2018. [Online]. Available: https://dl.acm.org/doi/10.5555/3235404

[4] Len Bass et al., "DevOps: A Software Architect's Perspective," Addison-Wesley Professional. [Online]. Available: https://ptgmedia.pearsoncmg.com/images/9780134049847/samplepages/9780134049847.pdf

[5] Mike Olumide, "A Comparative Analysis of Traditional vs. Generative AI-Powered Security Tools in Cloud Platforms," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/391346879_A_Comparative_Analysis_of_Traditional_vs_Generative_AI-Powered_Security_Tools_in_Cloud_Platforms

[6] Gaurav Agarwal, "AIOps Explained: Stages, Benefits and Use Cases," Hexaware Blog, 2025. [Online]. Available: https://hexaware.com/blogs/aiops-explained-stages-benefits-and-use-cases/

[7] Yinfang Chen et al., "Automatic Root Cause Analysis via Large Language Models for Cloud Incidents," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/379998668_Automatic_Root_Cause_Analysis_via_Large_Language_Models_for_Cloud_Incidents

[8] Pankaj Prasad, Padraig Byrne, and Gregg Siegfried, "Market Guide for AIOps Platforms," Gartner Research, 2022. [Online]. Available: https://www.gartner.com/en/documents/4015085

[9] Google, "2021 Accelerate State of DevOps Report,". [Online]. Available: https://services.google.com/fh/files/misc/state-of-devops-2021.pdf

[10] Sam Suthar, "What is AIOps? A Complete Guide to AI-Powered IT Operations," Middleware Blog, 2025. [Online]. Available: https://middleware.io/blog/what-is-ai-ops/

Journal of Information Systems Engineering and Management
2025, 10(61s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

[11] Jeffrey O. Kephart and D.M. Chess, "The Vision Of Autonomic Computing," ResearchGate, 2003. [Online]. Available: https://www.researchgate.net/publication/2955831_The_Vision_Of_Autonomic_Computing

[12] Tim Fountaine, Brian McCarthy, and Tamim Saleh, "Building the AI-Powered Organization," Harvard Business Review, 2019. [Online]. Available: https://hbr.org/2019/07/building-the-ai-powered-organization