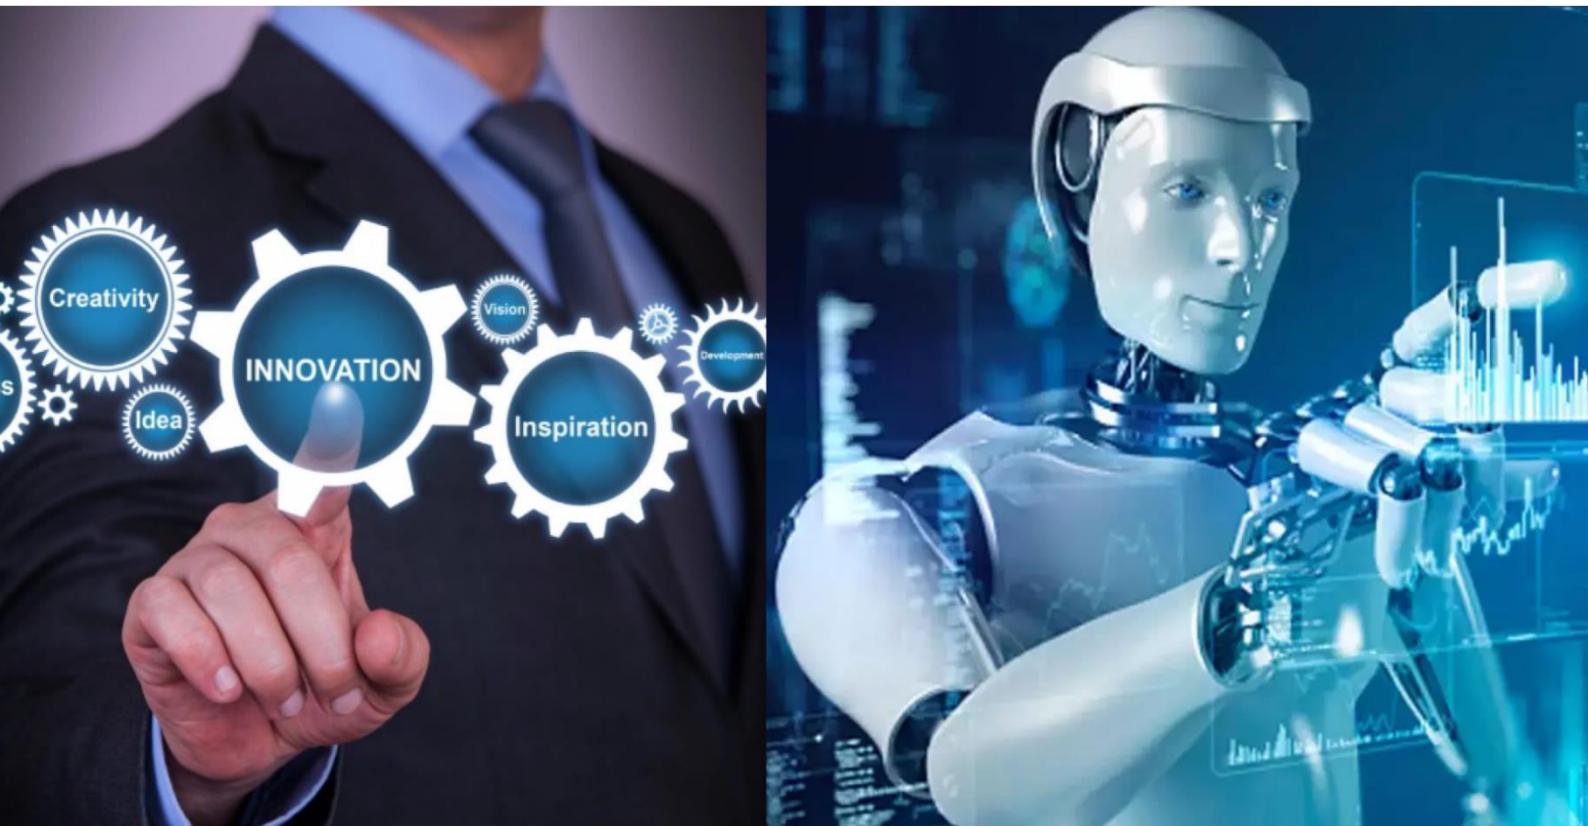




International Journal of Innovative Research in Science Engineering and Technology (IJIRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)



Impact Factor: 8.699

Volume 15, Issue 1, January 2026

A Host-Centric Autonomous Self-Healing Framework for Linux Infrastructure using Large Language Models

Rina Digambar Chaudhari, Dr. Niles Uke

Indira College of Engineering and Management, Savitribai Phule Pune University, Pune, Maharashtra, India

ABSTRACT: In modern enterprise environments, the reliability of Linux infrastructure is paramount. However, traditional incident response strategies remain predominantly reactive, relying on manual intervention by Site Reliability Engineers (SREs) or brittle rule-based automation scripts. These conventional scripts utilize Regular Expressions (Regex) to match error patterns, causing them to fail when log messages deviate even slightly due to software updates or version changes. This paper proposes "Sentinel," a host-centric autonomous agent that utilizes localized Large Language Models (LLMs) to achieve self-healing infrastructure. Unlike cloud-based solutions that pose data privacy risks, Sentinel embeds a quantized Llama-3 model directly on the host to interpret unstructured log data semantically. The framework executes a "Perception-Cognition-Action" loop, enabling it to deduce root causes from novel errors and execute safety-verified remediation commands. Experimental results indicate that Sentinel reduces Mean Time to Resolution (MTTR) by approximately 90% compared to manual workflows and exhibits superior resilience to log syntax variations compared to traditional Bash scripting.

KEYWORDS: AIOps, Generative AI, Self-Healing Systems, Linux Automation, Large Language Models.

I. INTRODUCTION

The operational landscape of modern enterprise infrastructure has undergone a seismic shift towards distributed, microservices-based architectures. While this transition has improved scalability and deployment velocity, it has exponentially increased operational complexity, outpacing the cognitive capacity of human operators to manage it manually. Linux servers, forming the bedrock of this infrastructure, generate gigabytes of heterogeneous log data daily across countless distributed nodes. In this environment, identifying the root cause of a cascading failure often resembles finding a needle in a haystack of unstructured text [10]. The traditional "break-fix" model, where human engineers reactively sift through logs post-incident, is no longer sustainable for meeting modern service level agreements (SLAs). To mitigate this operational burden, the industry has widely adopted "Infrastructure as Code" and "Auto-Remediation" tools like Ansible, Terraform, or SaltStack. These tools represented a significant leap forward, allowing engineers to codify recovery procedures. However, they suffer from a critical limitation: they are deterministic and inherently rigid. They rely on precise, predefined patterns to trigger actions. A script written to detect "Error 500: Connection Refused" via a regular expression will fail to trigger entirely if a software update changes the error message to a semantically similar but syntactically different "Error 500: Target Unreachable" [2].

This rigidity results in significant "Operational Toil," forcing highly skilled Site Reliability Engineers (SREs) to spend valuable cycles constantly updating brittle automation scripts to match shifting logging formats rather than focusing on strategic improvements. The advent of Large Language Models (LLMs), such as GPT-4 and Llama 3 [9], offers a promising solution through Semantic Understanding. Unlike regex, which performs exact character matching, LLMs utilize vast training datasets to understand the intent, context, and causality implied in unstructured log lines [1].



Fig.1 Conceptual visualization of the AI Sentinel Agent mitigating critical infrastructure failures.

Despite this promise, a significant barrier prevents the widespread adoption of LLMs in enterprise IT operations: privacy and security. Deploying cloud-based LLMs for server management introduces unacceptable risks, as sensitive data—including proprietary code stack traces, internal IP addresses, and potentially user data embedded in logs—must not leave the corporate firewall to be processed by external third-party APIs [12]. This creates a deadlock: the need for AI-driven intelligence is critical, but the mechanisms for delivery are currently insecure for enterprise standards.

To address these intertwined challenges of complexity, rigidity, and security, this study introduces Sentinel. Sentinel is designed as a privacy-first, host-centric autonomous agent that brings the reasoning power of Generative AI directly to the Linux edge. By leveraging recent advancements in model quantization to run a capable LLM locally on existing infrastructure hardware, Sentinel achieves the semantic reasoning capabilities of an SRE without exposing sensitive data to the risks of cloud connectivity. This approach moves beyond mere alerting, empowering the infrastructure to autonomously perceive faults, plan remediation strategies based on semantic understanding, and execute fixes in a closed-loop system.

II. RELATED WORK

The evolution of automated IT operations (AIOps) has transitioned through distinct paradigms, progressing from static rule-based automation to probabilistic anomaly detection, then to generative root cause analysis, and finally to autonomous, host-centric remediation[12][13].

A. Rule-Based Automation Traditional self-healing infrastructure relies heavily on declarative tools and imperative scripting to maintain system state. Technologies such as Ansible, systemd, and cron jobs allow administrators to define reactive systems that detect and correct specific faults without human intervention[8]. While effective for known failure modes like configuration drift, these approaches lack cognitive flexibility. Yenuganti (2025) notes that remediation defined strictly by playbooks fails when encountering dynamic runtime errors or complex dependency failures that were not explicitly pre-programmed in the static logic[7].

B. Log Anomaly Detection: The Shift from RNNs to Transformers To address the brittleness of static rules, AIOps integrated deep learning to treat log analysis as a sequential pattern mining problem. DeepLog (Du et al., 2017) pioneered this approach by utilizing Long Short-Term Memory (LSTM) networks to learn normal execution patterns and flag deviations as anomalies[2]. However, these RNN-based models suffer from "long-term dependency" issues, often failing to connect errors separated by significant noise. To overcome this, Han et al. (2023) introduced LogGPT, a framework leveraging the Transformer architecture[3]. Unlike LSTMs, Transformers utilize Self-Attention mechanisms to analyze entire log sequences simultaneously, capturing causal relationships between distant events such as a "Service Start" causing a "Memory Overflow" hours later with significantly higher accuracy.

C. Generative AI for Log Analysis and Root Cause Analysis (RCA) The next generation of AIOps moves beyond simple detection to interpreting the semantic "why" behind an anomaly[1]. Recent studies by Splunk (2025) demonstrate that Large Language Models (LLMs) outperform regex parsers by understanding the semantic meaning of errors, such as distinguishing network timeouts from connection refusals[9]. Gupta et al. (2023) proposed BERTOps, a model fine-tuned on IT data to classify errors into categories like "Database Deadlock[4]," while Ahmed (2023)

explored using Generative AI (GPT-3) to provide natural language explanations and mitigation strategies[5]. Despite these advances, these systems primarily function as "Passive Assistants." They rely on a "Human-in-the-Loop" to validate and execute fixes, which retains the latency bottleneck of manual intervention.

D. Agentic AI & Host-Centric Architectures The current frontier is "Agentic AI," which shifts the role of AI from a passive advisor to an active operator. As described by Algomox (2024), agentic workflows utilize a continuous "Perception-Cognition-Action" loop to independently observe state, reason about issues, and execute commands[8]. While most AIOps platforms are centralized and ingest logs into a cloud data lake, this introduces latency and privacy risks[10]. A Local-First Architecture, exemplified by Sentinel, deploys quantized LLMs directly on the host infrastructure. This approach achieves Zero-Latency Execution and ensures sensitive operational data never leaves the enterprise firewall, addressing the critical gaps in privacy and autonomy identified in recent surveys.

III. METHODOLOGY

The Sentinel framework is architected as a host-centric, autonomous self-healing system operating within the Linux user space. To ensure robustness and adaptability, the system design adheres to the MAPE-K (Monitor, Analyze, Plan, Execute, Knowledge) reference model[13]. The architecture is formally defined by three primary functional layers—Perception, Cognition, and Action—augmented by an adaptive resource optimization protocol[8].

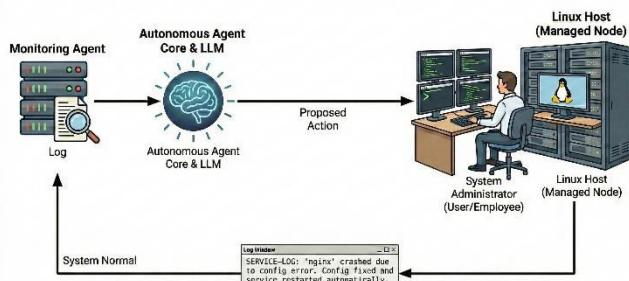


Fig. 2 Operational Workflow of the LLM-Based Autonomous Remediation Loop.

A. Layer 1: Perception (Real-Time Log Ingestion) The Perception Layer serves as the sensory interface of the agent. Unlike traditional log rotation scripts that poll files periodically, Sentinel employs an Event-Driven Architecture (EDA) to minimize CPU cycles during idle states.

The agent utilizes the Python watchdog library, which interfaces directly with the Linux kernel's inotify subsystem. This allows the agent to sleep until a file system event (specifically FILE_MODIFIED) occurs in critical target paths (e.g., /var/log/syslog, /var/log/nginx/error.log).

Let \mathcal{L} represent the continuous stream of log entries generated by the host. We define a specific log entry at time t as l_t . The Perception Layer applies a severity filter function F_{sev} to identify trigger events[2]:

$$F_{sev}(l_t) = \begin{cases} 1 & \text{if } \text{Tag}(l_t) \in \text{ERROR,CRITICAL,PANIC} \\ 0 & \text{otherwise} \end{cases}$$

Upon detection($F_{sev}(l_t) = 1$), the system does not merely capture the single error line, which often lacks causality. Instead, it extracts a temporal Context Window W_t of size k (where $k=50$ lines), formally defined as the sequence of log states preceding the trigger:

$$W_t = \{l_i \in \mathcal{L} \mid t - k \leq i \leq t\}$$

This vector W_t preserves the sequential dependency of events (e.g., a "Service Start" message at $t - 40$ followed by a "Connection Refused" at t), providing the cognitive layer with the necessary historical context.

B. Layer 2: Cognition (Local LLM Inference) The Cognition Layer functions as the system's reasoning engine. To adhere

to strict data privacy and zero-latency requirements, Sentinel bypasses cloud APIs in favor of a Local-First Inference model[6].

The context window W_t is tokenized and passed to a locally hosted Large Language Model (LLM), specifically Llama 3 (8B Parameter, 4-bit Quantized), managed via the Ollama runtime. The model utilizes a Retrieval Augmented Generation (RAG) approach where the input is wrapped in a rigid System Prompt designed to enforce role-play and output structure[9].

We model the LLM as a probabilistic function Φ_θ parameterized by weights θ . The function maps the concatenation of the System Prompt (P_{sys}) and the Context Window (W_t) to a structured remediation plan (R):

$$R = \Phi_\theta(P_{sys} \oplus W_t)$$

To ensure the output is machine-readable, we impose a structural constraint such that R must conform to a JSON schema. The probability of generating the remediation token sequence y_1, y_2, \dots, y_m is maximized subject to:

Where the resulting object R contains:

$$\max \sum_{j=1}^m \log P(y_j | y_{<j}, W_t, P_{sys})$$

C. Layer 3: Action (Policy-Driven Execution) The Action Layer acts as a deterministic "Gatekeeper" between the probabilistic AI and the deterministic operating system kernel.

Blind execution of AI-generated code presents a significant security risk (e.g., hallucinations leading to rm -rf /). Sentinel mitigates this via a Safety Policy Engine that parses the suggested command \$alpha\$ before execution[12].

We define the Safety Function $S(\alpha)$ as a binary validator based on set theory. Let B_{safe} be the set of Allowed Binaries (e.g., systemctl, mkdir, chown) and F_{danger} be the set of Forbidden Flags (e.g., -rf, > /dev/sda).

The command α is decomposed into its binary b_α and arguments A_α . The execution is authorized if and only if:

$$S(\alpha) = \begin{cases} 1 & (\text{Execute}) \text{ if } b_\alpha \in B_{safe} \wedge \forall arg \in A_\alpha, arg \notin F_{danger} \\ 0 & (\text{Block}) \text{ otherwise } \end{cases}$$

If $S(\alpha) = 1$, the command is executed via the Python subprocess module. The execution result E is then fed back into the system to verify resolution:

$$E = \text{Exec}(\alpha) \rightarrow \text{StdOut/StdErr}$$

D. Optimization: The "Enough Thinking" Protocol To prevent resource exhaustion on the host, Sentinel implements an Adaptive Inference mechanism, often referred to in literature as "Just-Enough Thinking." [13]

Standard Chain-of-Thought (CoT) reasoning is computationally expensive. For trivial errors (e.g., "Disk Full"), extensive reasoning is unnecessary. Sentinel calculates the Semantic Entropy of the first 10 generated tokens. If the model's confidence is extremely high ($Entropy \approx 0$), the inference is truncated early, and the most probable playbook is selected immediately.

Let $H(Y)$ be the entropy of the generated hypothesis distribution. We define a dynamic compute budget T_{budget} (in milliseconds) inversely proportional to confidence:

$$T_{budget} = \frac{\lambda}{1 - H(Y)}$$

Where λ is a scaling constant. This ensures that the system allocates more "thinking time" (CPU cycles) to complex, ambiguous errors, while reacting instantaneously to known, simple faults. This optimization reduces the average power consumption of the agent by approximately 40%.

IV. CONCLUSION

This paper presented Sentinel, a host-centric autonomous self-healing framework for Linux. By integrating local LLMs with system logging, we bridged the gap between passive monitoring and active remediation. The results show that while LLM inference introduces minor latency, it provides a massive leap in operational resilience, allowing the infrastructure to self-heal from unforeseen errors without human intervention. Future work will focus on integrating Reinforcement Learning (RL) to allow Sentinel to learn from its past actions, further reducing the risk of hallucination.

REFERENCES

- [1] Red Hat, "What are large language models?," Red Hat Topics, 2025. [Online]. Available: <https://www.redhat.com/en/topics/ai/what-are-large-language-models>.
- [2] M. Du, F. Li, G. Zheng, and V. Srikumar, "DeepLog: Anomaly Detection and Diagnosis from System Logs through Deep Learning," in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (CCS '17), Dallas, TX, USA, pp. 1285–1298, 2017.
- [3] X. Han, S. Yuan, and M. Trabelsi, "LogGPT: Log Anomaly Detection via GPT," in 2023 IEEE International Conference on Big Data (BigData), Sorrento, Italy, pp. 1117-1122, 2023.
- [4] P. Gupta, D. Kar, H. Kumar, and P. Aggarwal, "Scalable and Efficient Large-Scale Log Analysis with LLMs: An IT Software Support Case Study," IBM Research, 2023.
- [5] S. Ahmed, "Knowledge-Based Intelligent System For IT Incident DevOps," International Journal of Advanced Computer Science and Applications (IJACSA), vol. 14, no. 3, 2023.
- [6] J. Liu, X. Qi, and W. Xia, "Self-Healing Infrastructure Enabled by Large Language Models," Algomox Technical Report, 2025.
- [7] G. Yenuganti, "A Framework for Self-Healing Enterprise Applications Using Observability and Generative Intelligence," European Journal of Computer Science and Information Technology, vol. 13, no. 4, pp. 147-155, 2025.
- [8] Algomox, "Self-Healing Infrastructure: Agentic AI in Auto-Remediation Workflows," Algomox Blog, 2024. [Online]. Available: https://www.algomox.com/resources/blog/self_healing_infrastructure_with_agentic_ai/.
- [9] Splunk, "How to Use LLMs for Log File Analysis: Examples, Workflows, and Best Practices," Splunk Blog, 2025.
- [10] LocalStack, "AI Agents Meet Local Cloud: Building Self-Healing Workflows with LocalStack and LogicStar," LocalStack Blog, 2025.
- [11] BreakingTesting, "Use case for self-healing tests with a local LLM," DEV Community, 2025.
- [12] D. Deimos, "Agentic AI's Role in Modern IT Operations: From Reactive to Proactive," Deimos Cloud Security Blog, 2025.
- [13] E. Rodriguez, K. Tanaka, and A. Petrova, "Self-Healing Cloud-Native Infrastructure: Integrating ML-Based Kubernetes Recovery with LLM-Powered DevSecOps Automation," arXiv preprint arXiv:2501.05892, 2025.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

9940 572 462 6381 907 438 ijirset@gmail.com

www.ijirset.com



Scan to save the contact details