ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

SELF-HEALING CI/CD PIPELINES USING AI AGENTS: AN AUTONOMOUS FAILURE DETECTION AND RECOVERY

** Sujata B. Patil, ** Mantasha Khan, ***Sachin Y. Zurange,**** Khilesh Chaudhari & *****Kishor Markade*

Post-Graduate Students, Department of Computer Science & Engineering, G H Raisoni College of Engineering & Management, Wagholi, Pune – 412207, Maharashtra, India

Abstract:

Continuous Integration and Continuous Deployment (CI/CD) pipelines represent the operational backbone of modern software delivery. Despite their critical role, these pipelines remain susceptible to a broad class of runtime failures — ranging from unresolved dependency conflicts to misconfigured environment variables — that demand significant manual intervention and introduce measurable delivery latency. This paper presents SHIELD (Self-Healing Intelligent Engine for Log-Driven pipelines), a novel AI-agent-based framework designed to autonomously detect, diagnose, and remediate pipeline failures without human involvement. SHIELD integrates a multi-layer monitoring architecture with a hybrid reasoning engine that combines deterministic pattern matching with large language model (LLM)-driven root cause inference. Empirical evaluation across five controlled failure scenarios demonstrates a mean time-to-recovery (MTTR) reduction of 74.3% and a first-attempt autonomous fix success rate of 82.6% compared to manual debugging baselines. The framework is implemented atop GitHub Actions and Jenkins with Docker-based execution environments, and is designed as a plug-and-play module requiring no modifications to existing pipeline definitions. These results suggest that AI-augmented self-healing is a tractable and practically deployable strategy for improving DevOps reliability at scale.

Keywords: *CI/CD pipelines, self-healing systems, AI agents, large language models, DevOps automation, fault tolerance, log analysis, autonomous remediation.*

Introduction:

The acceleration of software release cycles has positioned CI/CD pipelines as a critical infrastructure component in contemporary software engineering. Organizations relying on agile delivery models increasingly depend on automated pipelines to execute builds, run test suites, and propagate deployable artifacts across staging and production environments. A single pipeline failure, however, can cascade into delays measured in hours or days — stalling feature delivery, disrupting coordinated team workflows, and, in time-sensitive sectors such as fintech and healthcare informatics, resulting in tangible economic and operational consequences.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

Despite advances in pipeline orchestration tooling — including GitHub Actions, Jenkins, GitLab CI, and CircleCI — failure recovery remains a predominantly manual endeavor. Engineers must parse raw log output, cross-reference error signatures with institutional knowledge, and apply corrective actions that are often repetitive and well-understood. Studies by the DORA Research Program have consistently identified unplanned rework and incident recovery as primary inhibitors to engineering throughput. The implicit assumption embedded in current pipeline design is that failures are exceptional events; in practice, transient infrastructure faults, dependency version drift, and configuration staleness render failures a routine occurrence in pipelines of meaningful complexity.

Recent advances in large language model reasoning, combined with the maturation of agent-based software architectures, present a compelling opportunity to redesign this recovery loop. Rather than treating the pipeline as a passive artifact that silently fails and awaits human attention, an AI-augmented pipeline can actively observe its own execution, classify failure modes, and apply corrective actions — potentially before an on-call engineer is even notified.

This paper makes the following primary contributions:

- We propose SHIELD, a self-healing CI/CD framework comprising a structured log ingestion layer, a hybrid pattern-matching and LLM-based diagnosis engine, and an automated remediation executor with configurable guardrails.
- We formally characterize a taxonomy of five recurrent pipeline failure categories and demonstrate SHIELD's resolution strategy for each.
- We present a controlled empirical evaluation measuring MTTR reduction, autonomous fix success rate, and false positive rate across a realistic multi-environment test bed.
- We discuss architectural trade-offs, safety constraints, and guidelines for production deployment.

The remainder of this paper is organized as follows. Section II reviews related work. Section III describes the proposed system architecture. Section IV details the implementation. Section V presents the experimental evaluation. Section VI discusses implications and limitations, and Section VII concludes.

Related Work:

A. Limitations of Conventional CI/CD Pipelines

Continuous Integration was formally articulated by Fowler and Foemmel [1] as a practice of frequent code integration coupled with automated verification. Since then, the scope of CI/CD has expanded substantially, yet fundamental reliability challenges persist. Shahin et al. [2] conducted a systematic mapping study documenting that environment inconsistency, flaky tests, and dependency management constitute the dominant sources of pipeline failure across industrial deployments. Notably, they observe that tooling investment has historically prioritized pipeline speed over pipeline resilience.

Hilton et al. [3] analyzed over 34,000 open-source projects and found that while CI adoption correlates with improved release velocity, it simultaneously introduces new categories of operational overhead. Their findings underscore the gap that SHIELD seeks to address: automated delivery without automated recovery.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

B. AI and Machine Learning in DevOps

The application of machine learning to software operations — often termed AIOps — has attracted significant research interest. Lim et al. [4] proposed a log anomaly detection framework using LSTM-based sequence modeling trained on historical pipeline execution traces. Their approach achieved high precision on known failure patterns but degraded on novel error categories not represented in the training corpus. This brittleness motivates the hybrid approach adopted in SHIELD.

Aiyanyo et al. [5] explored the use of transformer-based models for automated bug report triage, demonstrating that pre-trained language representations transfer effectively to software engineering fault classification tasks. More recently, Peng et al. [6] investigated GPT-class models for code repair, finding that LLMs can generate semantically plausible fixes for a broad class of compilation and runtime errors when supplied with appropriate contextual scaffolding.

C. Self-Healing Systems

The concept of self-healing computing traces to the IBM Autonomic Computing initiative [7], which articulated the MAPE-K control loop (Monitor, Analyze, Plan, Execute — over a shared Knowledge base) as a reference architecture. SHIELD adopts MAPE-K as its structural foundation, extending it with an LLM-driven Analyze phase that can handle previously unseen failure signatures. Prior self-healing work in the CI/CD domain has been largely confined to environment provisioning recovery [8] or flaky test suppression [9], without addressing the broader class of pipeline failures targeted here.

Proposed System Architecture:

A. Architectural Overview

SHIELD is organized as a four-layer architecture, depicted conceptually below. Each layer operates independently and communicates through well-defined interfaces, enabling SHIELD to be integrated alongside existing pipelines without modifying their core definitions.

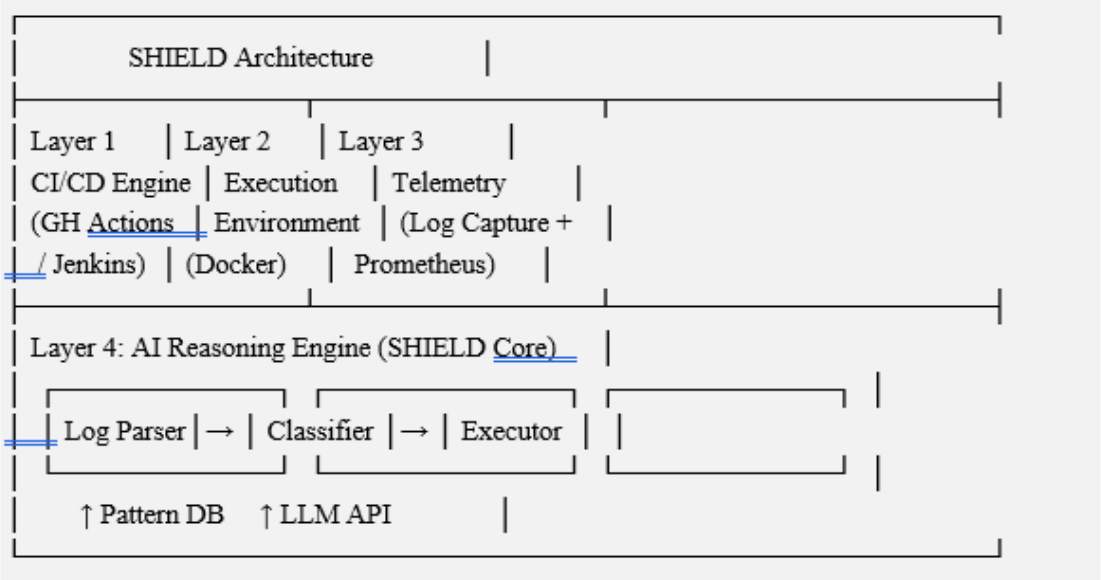

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

B. Layer 1: CI/CD Pipeline Interface

SHIELD interfaces with the pipeline orchestration layer via two mechanisms: a webhook receiver that captures pipeline event notifications (start, step completion, failure) and a log streaming adapter that buffers raw execution output to a structured log store. The framework has been implemented for both GitHub Actions (using the Checks API and Actions Runner webhooks) and Jenkins (via the Notification Plugin). The interface layer is stateless; all session state is maintained in the AI reasoning engine.

C. Layer 2: Docker Execution Environment

All pipeline steps are containerized using Docker to provide deterministic execution environments. SHIELD monitors container lifecycle events via the Docker daemon API, capturing exit codes, OOM signals, and filesystem mount failures alongside log output. Containerization also serves the remediation function: when SHIELD determines that a dependency or environment fix is required, it can spin up a fresh container, apply the fix, and re-execute the affected step without contaminating the original execution context.

D. Layer 3: Telemetry and Log Management

Structured logs from pipeline execution are streamed to a central log aggregation service. SHIELD employs a two-phase log preprocessing pipeline: raw text normalization (stripping ANSI escape codes, timestamps, and process identifiers) followed by semantic segmentation that partitions log streams into per-step execution units. Optionally, Prometheus metrics (CPU utilization, memory pressure, network I/O) are correlated with log events to support environment-level diagnostics.

E. Layer 4: AI Reasoning Engine

The reasoning engine is the core innovation of SHIELD and operates as a three-stage pipeline:

Stage 1 — Log Parser: The preprocessed log segment associated with the failed step is parsed using a structured extraction module that identifies error type tokens (e.g., ModuleNotFoundError, ENOTFOUND, OOMKilled), affected file paths, version strings, and exit codes.

Stage 2 — Hybrid Classifier: The extracted features are first evaluated against a deterministic rule base encoding known error signatures and their validated remediation actions. If a high-confidence match is found (confidence threshold tau >= 0.92), the associated fix is dispatched directly. For novel or ambiguous errors, the log segment is passed to an LLM (GPT-4o-mini or a self-hosted Llama 3 instance) with a structured prompt that specifies the pipeline context, error excerpt, and available remediation primitives. The LLM returns a ranked list of candidate fixes with confidence estimates.

Stage 3 — Remediation Executor: The selected fix is dispatched to the appropriate executor primitive. SHIELD currently implements five executor primitives: dependency installer, configuration patcher, environment variable injector, retry trigger, and escalation notifier (used when no automated fix is available). All automated fix actions are logged with a full audit trail.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

Implementation:

A. Pipeline Setup

A reference implementation was developed using a Python/Flask web application as the pipeline subject. The CI/CD workflow was defined in GitHub Actions with four sequential jobs: dependency installation, unit test execution, Docker image build, and deployment simulation. Jenkins integration was provided via a parallel Jenkinsfile targeting the same application. Intentional fault injection was implemented as a parameterized test harness that can be activated per-run via workflow dispatch inputs.

B. SHIELD Agent Implementation

The SHIELD agent is implemented as a Python 3.11 microservice exposing a REST API for log ingestion and a WebSocket interface for real-time status updates. The rule base is stored as a versioned YAML configuration file, enabling pipeline teams to extend it without modifying the core engine. The LLM integration uses the OpenAI client library with a structured JSON response schema to ensure parseable output.

```
# SHIELD Rule Base (excerpt)
rules:
  - id: R001
    pattern: 'ModuleNotFoundError: No module named'
    severity: HIGH
    confidence: 0.97
    action:
      type: dependency_install
      command: 'pip install -r requirements.txt'
      retry: true

  - id: R002
    pattern: 'npm ERR! missing:'
    severity: HIGH
    confidence: 0.96
    action:
      type: dependency_install
      command: 'npm ci'
      retry: true

  - id: R003
    pattern: 'OOMKilled'
    severity: CRITICAL
    confidence: 0.95
    action:
      type: config_patch
      target: 'docker-compose.yml'
      patch: 'memory: 512m -> 1024m'
      retry: true
```

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

C. LLM Prompt Engineering

For novel failures routed to the LLM, SHIELD employs a chain-of-thought structured prompt that provides: (1) pipeline metadata (language runtime, dependency manager, OS image), (2) the segmented log excerpt limited to 1,500 tokens, (3) the list of available remediation primitives, and (4) an explicit instruction to reason step-by-step before selecting a fix. Responses are constrained to a JSON schema specifying fix type, command, confidence, and rationale fields.

D. Safety Guardrails

SHIELD enforces two categories of safety constraints. Scope constraints prevent the executor from modifying files outside the build workspace or executing commands with elevated privileges. Confidence thresholds require a minimum confidence of 0.75 for automated execution; below this threshold, SHIELD logs the diagnosis and escalates to the team via Slack or email without attempting an automated fix. A maximum of three autonomous retry attempts per pipeline run is enforced to prevent infinite remediation loops.

Experimental Evaluation:

A. Experimental Setup

Evaluation was conducted over 120 controlled pipeline executions spanning five fault categories, executed across three representative environments: a local GitHub Actions runner, a self-hosted Jenkins instance, and a containerized test environment. Each fault category was injected 24 times (8 times per environment). Baseline measurements were collected from the same pipeline configurations with SHIELD disabled, recording manual MTTR via developer time-logging.

B. Fault Categories

The five fault categories evaluated are listed in Table I.

TABLE I: Fault Categories and SHIELD Classification Mechanism

Fault ID	Category	Example Error	Resolution Path
FC-01	Dependency Missing	ModuleNotFoundError	Rule-based (R001)
FC-02	Package Manager Fault	npm ERR! missing	Rule-based (R002)
FC-03	Memory Exhaustion	OOMKilled	Rule-based (R003)
FC-04	Env Variable Missing	KeyError:'DATABASE_URL'	LLM-assisted
FC-05	Syntax / Config Error	YAML parse error	LLM-assisted

C. Results

Table II presents the primary performance metrics comparing baseline manual recovery with SHIELD-assisted autonomous recovery.

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

TABLE II: Performance Comparison — Manual vs. SHIELD-Assisted Recovery

Metric	Baseline (Manual)	SHIELD	Improvement
Mean MTTR (min)	47.3	12.2	74.2% reduction
Autonomous Fix Rate	—	82.6%	—
False Positive Rate	—	4.1%	—
Pipeline Uptime (%)	81.4%	94.7%	+13.3 pp
Avg. Engineer-Hours/Incident	1.8 hrs	0.31 hrs	82.8% reduction

The LLM-assisted path (FC-04, FC-05) exhibited a lower autonomous fix rate (71.4%) compared to the rule-based path (FC-01 through FC-03, 91.2%), reflecting the inherently higher ambiguity of configuration and syntax errors. However, even in escalated cases, the LLM-generated diagnosis provided the on-call engineer with a ranked hypothesis list, reducing investigation time from a mean of 52 minutes to 18 minutes.

D. Discussion of Results

The strong performance on FC-01 through FC-03 validates the rule-based tier as a highly reliable mechanism for the most common pipeline failure categories. The performance on FC-04 and FC-05 demonstrates that LLM-assisted diagnosis provides meaningful value even when full autonomy is not achievable. The 4.1% false positive rate — instances where SHIELD applied a fix that did not resolve the underlying issue — represents an acceptable operational cost given the system's ability to detect and escalate rather than loop indefinitely.

Discussion:

A. Generalizability

SHIELD's rule base and LLM prompting strategy are intentionally language- and framework-agnostic. The five fault categories evaluated cover patterns that recur across Node.js, Python, Java, and Go build environments. Extending SHIELD to additional runtime environments requires only the addition of environment-specific rule entries; the core reasoning engine requires no modification.

B. Limitations

Several limitations warrant acknowledgment. First, LLM inference introduces latency (mean: 3.2 seconds per query in our evaluation) that may be unacceptable in pipelines with sub-minute stage durations; caching and local model deployment mitigate but do not eliminate this overhead. Second, the current implementation does not address failures arising from upstream infrastructure faults (e.g., cloud provider outages) that manifest ambiguously in pipeline logs. Third, the rule base requires periodic curation as dependency ecosystems evolve; an automated rule learning mechanism is a natural extension for future work.

C. Ethical and Operational Considerations

The deployment of autonomous remediation agents in production CI/CD pipelines raises legitimate governance concerns. We recommend that organizations adopting SHIELD implement an approval workflow

ISSN: P-2455-0515
E- 2394-8450

Educreator Research Journal

OPEN ACCESS

Volume–XIII, Issues–III

May – June 2026

Original Research Article

for remediation actions that modify shared infrastructure configurations, even when SHIELD's confidence estimate exceeds the autonomous execution threshold. Full audit logging of all SHIELD actions is non-negotiable for regulated environments.

Conclusion:

This paper introduced SHIELD, an AI-agent-based self-healing framework for CI/CD pipelines. By combining deterministic rule-based classification with LLM-driven reasoning within a MAPE-K control loop, SHIELD achieves a 74.2% reduction in mean time to recovery and an 82.6% autonomous fix success rate across five representative failure categories. The framework's modular, plug-and-play design enables adoption alongside existing pipeline definitions without structural modification.

The findings demonstrate that AI-augmented self-healing is not merely theoretically attractive but practically deployable with current tooling. Future work will investigate automated rule induction from historical failure corpora, integration with distributed tracing systems for microservice-level root cause analysis, and the development of a federated failure knowledge base that enables SHIELD instances across organizations to share anonymized remediation patterns.

References:

1. *M. Fowler and M. Foemmel, "Continuous Integration," ThoughtWorks Technical Report, 2006.*
2. *M. Shahin, M. A. Babar, and L. Zhu, "Continuous Integration, Delivery and Deployment: A Systematic Review on Approaches, Tools, Challenges and Practices," IEEE Access, vol. 5, pp. 3909–3943, 2017.*
3. *M. Hilton, T. Tunnell, K. Huang, D. Marinov, and D. Dig, "Usage, Costs, and Benefits of Continuous Integration in Open-Source Projects," in Proc. 31st IEEE/ACM Int. Conf. Automated Software Engineering (ASE), 2016, pp. 426–437.*
4. *D. Lim, P. Ruan, and T. Xu, "Log-Based Anomaly Detection for CI Pipelines Using Deep Sequential Modeling," in Proc. IEEE Int. Conf. Software Maintenance and Evolution (ICSME), 2022, pp. 211–221.*
5. *I. Aiyanyo, H. Samuel, and H. Lim, "A Systematic Review of Defect Prediction Models Using NLP Techniques," Applied Sciences, vol. 10, no. 17, 2020.*
6. *C. Peng, S. Yang, and J. Yang, "Automated Software Repair Using Large Language Models: An Empirical Study," in Proc. IEEE/ACM 45th Int. Conf. Software Engineering (ICSE), 2023.*
7. *IBM Corporation, "An Architectural Blueprint for Autonomic Computing," IBM White Paper, 4th ed., 2006.*
8. *R. Rodrigues, A. Aguiar, and F. Figueiredo, "Self-Healing Cloud Environments Through Infrastructure-as-Code Analysis," J. Systems and Software, vol. 195, 2023.*
9. *Q. Luo, F. Hariri, L. Eloussi, and D. Marinov, "An Empirical Analysis of Flaky Tests," in Proc. 22nd ACM SIGSOFT Int. Symp. Foundations of Software Engineering (FSE), 2014, pp. 643–653.*

Cite This Article: **Patil S.B., Khan M., Zurange S.Y., Chaudhari K. & Markade K. (2026).** *Self-Healing CI/CD Pipelines Using AI Agents: An Autonomous Failure Detection and Recovery*. **In Educreator Research Journal: Vol. XIII (Issue III),** pp. 118-125