



Original Article

LLM-Based Auto-Remediation Model for DevOps Pipeline Failures

Siva Sai Krishna Suryadevara
Sr. AEM Developer at Maganti IT Resources, USA.

Abstract - DevOps pipelines are the backbone of modern software delivery. They make it very easy to quickly, automatically, and continuously integrate and deploy across many other different cloud environments. However, these pipelines often fail because of configuration drifts, dependency mismatches, infrastructure inconsistencies as well as code-level errors. This causes operational bottlenecks that slow down releases along with their require a lot of human intervention. Standard remediation techniques rely heavily on these pre-written scripts, rigid rule-based frameworks, or engineers doing manual diagnostics. These methods often don't work well with the changing nature of toolchains and multi-cloud architectures. Recent progress in Large Language Models (LLMs) could make these pipelines more reliable by using context-aware reasoning, advanced root-cause analysis, and automated solution formulation. This work presents an LLM-based auto-remediation methodology designed to detect pipeline failures, analyze logs, correlate error patterns, provide corrective actions, and when feasible autonomously rectify issues in actual time. The proposed paradigm integrates natural language understanding with DevOps telemetry, source code analysis & configuration verification to connect detection as well as resolution. Experimental evaluations of widely used these CI/CD systems demonstrate substantial improvements in mean time to recovery (MTTR), a reduction in repetitive manual debugging tasks, and enhanced their pipeline dependability, especially in complex cloud-native workflows. The LLM-driven system is better than typical remediation scripts because it can adapt to the latest types of failures without needing to be reprogrammed, provide better information about the context of mistakes, and works well across teams & environments. This study shows how LLMs can change DevOps operations from being reactive to proactive, cut down on interruptions & let engineers focus on coming up with the latest ideas instead of putting out fires. This is a huge step toward self-healing automation in modern software delivery ecosystems.

Keywords - DevOps, CI/CD, Auto-Remediation, Large Language Models, AI-Ops, Pipeline Failure Detection, Root Cause Analysis, Automation, Observability, Intelligent Troubleshooting.

1. Introduction

Over the past ten years, the way modern software is delivered has changed at an amazing rate. Companies are using cloud-native designs more and more, adding the latest features faster, and relying heavily on their automation to keep things stable at scale. The DevOps pipeline is the heart of this change. It is a continuous integration and continuous delivery (CI/CD) system that brings together code, infrastructure, testing, security as well as deployment into one smooth process. The first simple way to design and deploy applications has quickly grown into a complex system of tools, microservices, cloud platforms & infrastructure-as-code (IaC) operations. DevOps has greatly sped up release cycles, but it has also created a whole new set of these operational problems, especially when it comes to finding, fixing, and preventing failures.

1.1. Background Development of DevOps Pipelines

At first, DevOps pipelines were very simple: write code, run basic tests, and deploy to one environment. Software architectures have changed from monolithic apps to distributed microservices over time. Companies used multi-cloud tactics to save money, improve reliability & reach customers all over the world. Infrastructure-as-a-ServiceCode became the norm, making it easier to declare these cloud resources and automatically manage environments.

This change brought about several other benefits, including scalability, modularity, and faster deployments. However, it also made the operational surface area much larger. A modern DevOps pipeline manages a lot of different steps that are all well connected to one other. These steps include building Docker images, deploying Kubernetes, running Terraform, checking configurations, running security scans, running integration tests, setting up monitoring hooks, and more. There is a danger of failure at every other step, with every toolkit and dependency. As pipelines grow, the chances of runtime failures, misconfigurations, dependency conflicts & environmental drift also grow.

Increasing Complexity Microservices break up the functionality of an application into smaller parts, which speeds up feature development but makes the pipeline very less stable. Each service may have its own way of building, testing, setting up the environment, and running dependencies. In a multi-cloud context, deployments may include AWS, Azure, and GCP all at

the same time, with each having its own APIs, networking protocols & resource configurations. At the same time, Infrastructure as Code (IaC) tools like Terraform, Pulumi, and CloudFormation make these things even more complicated by adding state management, version control, remote backends, and policy enforcement.

This complicated ecosystem makes a lot of logs, metrics, metadata, and other goods along the way. When anything goes wrong, engineers have to sort through a lot of information to find the main cause, often with little time to spare. As a result, the industry is gradually looking for ways to automate not just detection but also smart remediation.

DevOps teams have done a great job of automating development, testing as well as deployment processes, but fixing problems has mostly stayed a human job. Most CI/CD systems can tell engineers when something goes wrong, but only a few can figure out what the failure means or come up with ways to fix it on their own. Standard rule-based remediation scripts work well for problems that are easy to forecast, but they don't work when surroundings change, dependencies change, or new patterns emerge. This difference has made auto-remediation the next huge thing in the rise of DevOps, thanks to improvements in machine learning and, more recently, Large Language Models (LLMs).

1.2. Challenges

Even though tools have come a long way, teams still have to deal with a lot of long-term problems when pipelines fail:

- Higher Mean Time to Recovery (MTTR): When CI/CD workflows fail, they often need to be looked into by hand. Developers need to look at logs, analyze stack traces, check configurations, and try many different solutions. This makes MTTR longer and messes up delivery timetables.
- Disruptions in sound and notification overload: Pipelines create a lot of log information. A single failed deployment can create logs from Kubernetes, cloud APIs, Infrastructure as Code tools, and application services. Alerts can be a hassle, especially when a lot of failures are caused by the same problem.
- Not Enough Understanding of Contextual Mistakes: Most tools use pattern recognition or exit codes to find errors, but they don't look at the context. Network settings, IAM policies, problems with pod autoscaling, or previous DNS information might all cause a "connection timeout." Without context, automated technologies can't tell the difference.
- Dependence on manual intervention: Human knowledge is very important to mature DevOps teams. Troubleshooting still relies on the knowledge, experience & intuition that comes from years of dealing with many problems. Because of this dependence, cleanup is unpredictable and very hard.
- Diagnostic Methodologies That Don't Fit: Different engineers approach problems in many other different ways, which leads to different levels of quality in the fixes. Some people might carefully look over the logs, some might do the task again, and yet others might use quick fixes without really understanding the problem. The lack of consistency increases operational risk.
- Challenge in Identifying Fundamental Causes in Distributed Systems: In microservices and multi-cloud systems, problems often come from how different parts of the system work together, not from one part of the system. Keeping an eye on these dependencies by hand takes a lot of time. A failing build could be caused by an incorrectly set environment variable in a downstream service or a Terraform state that has changed in another region.
- These problems highlight a clear difference: while detection systems are quite very good, automated reasoning and fixing problems aren't very good yet.

1.3. Problem Statement

Modern CI/CD systems are good at finding many problems, but not so good at fixing them. They work more like alarm systems than like things that fix themselves. A number of important limits define the problem area:

- They find problems but don't fix them. Most pipeline solutions let teams know when something is wrong, but they can't take action to fix it unless they have pre-written scripts.
- In cloud systems that change all the time, static rules don't work. The infrastructure of the cloud changes too quickly. Static rules can't handle changing their dependencies, updates, or new patterns of errors.
- There isn't a single framework that makes it easy to think about logs, metrics, configuration files & deployment events. Current tools work on their own.
- Troubleshooting requires human intuition. Without better contextual intelligence, machines can't figure out why things happen or make accurate predictions about how to fix them.
- So, there is an urgent need for a dynamic, context-sensitive, self-learning system that thoroughly looks at problems and automatically takes steps to fix them.

1.4. Motivation

As businesses look to automate tasks on a broad scale, AIOps is becoming more popular. It uses machine learning to find, predict & improve anomalies. Because LLMs are now available, AIOps can now be used for tasks that require a lot of reasoning, such as figuring out what logs mean and finding the root cause of a problem.

1.4.1. LLM can do things like log reasoning, pattern recognition as well as code comprehension.

Large Language Models have shown that they can understand more complex logs very well.

- Recognizing mistake signatures in a number of tools
- Making suggestions for changes to the code
- Understanding Infrastructure as Code & YAML settings
- Finding patterns in systems that are spread out
- These skills make them the best choice for running next-generation auto-remediation systems.

1.4.2. Reducing Downtime and Operational Burden

Auto-remediation makes it more easier to find, diagnose & fix problems in actual time, which lowers MTTR. This cuts down on downtime & the stress on DevOps teams, so they can focus on coming up with the latest ideas instead of dealing with many problems.

1.4.3. Making Developers More Productive

Developers spend a lot of effort figuring out what is wrong with CI/CD. An advanced remediation method reduces this friction, allowing engineers to stay focused & productive.

1.4.4. Immediate Action to Reduce Business Impact

Failures that continue for hours or days without being fixed can make these deployments harder, push back product launches, and make the user experience worse. Real-time remediation makes sure that things stay the same & go smoothly in dynamic delivery situations.

2. Literature Review

DevOps teams have had to rely heavily on their automation, monitoring & advanced decision-making systems since software delivery methods have changed so quickly. As pipelines become more complicated, the industry has begun using AIOps and machine-learning-based automation to make deployment processes better and lighten the load on their operations. This literature review summarizes the main research fields that affect modern auto-remediation systems, such as AIOps frameworks, log analytics & LLM-driven debugging. It also points out the gaps in current research.

2.1. Summary of AIOps Frameworks

Gartner came up with the phrase "AIOps" to describe platforms that combine big-data analytics, machine learning along with their automation to make IT operations better. The first AIOps frameworks focused on their log aggregation, event correlation, and alert minimization. This made it easier for operations teams to spot more trends in huge amounts of data. These solutions usually combine dashboards for monitoring, tracing, and metrics with statistical models to find unusual activity in these applications or infrastructure.

IBM's Watson AIOps, Dynatrace DAVIS, and ServiceNow AIOps are all advanced frameworks that combine behavior profiling, causal analysis & predictive insights. They use time-series forecasting, clustering, and correlation engines to find strange system states before they break. Studies have examined the amalgamation of AIOps with IT Service Management (ITSM) to link incidents, change records, and problem information, hence improving root-cause analysis. However, these systems still rely on their static rules or pre-trained machine learning models that don't have deep contextual reasoning, especially when it comes to DevOps pipelines that are always changing.

2.2. Ways to Analyze Logs and Find Anomalies

You need log analytics to find out what's wrong with your pipeline. Standard log-based anomaly detection uses statistical thresholds, regular expression analysis & simple clustering techniques to find unusual patterns. Recent studies have focused on unsupervised learning techniques, such as Isolation Forest, Autoencoders, LSTM-based models & graph neural networks for the representation of log sequences.

Log parsing techniques like Drain, Spell, and LogCluster make it easier to turn semi-structured logs into structured templates that can then be used to find these anomalies. At the same time, sequence-based anomaly detection models try to find changes in time between log events. This helps themselves detect unusual patterns that don't fit into the typical flow of a pipeline.

Even while those enhancements have made the process simpler to discover problems, there continue to be issues in DevOps settings that have developed quickly, where logs evolve constantly, application transmission lines change regularly, and refresher courses models are extremely expensive. Also, log abnormalities might be indicative to many prospective root causes, thus more investigation will be required to figure out exactly what the right measures to take are.

2.3. Machine Learning-Based Root Cause Analysis Approaches

Recent studies integrate learning-based causation with observability signals (logs, metrics, traces) to ascertain the most probable cause of a failure. Graph-based RCA models show how services are connected & find bottlenecks in these distributed systems. Attention-based deep learning models, on the other hand, focus on the signals that are most important to an event. Some researchers have tried to improve their accuracy by cross-correlating measurements and logs.

Still, ML-driven RCA has two problems: (1) it relies heavily on their information, therefore it needs large, labeled datasets of past failures, and (2) it has trouble adapting to the latest technologies or pipeline topologies without having to be retrained. Additionally, even after pinpointing the core issue, these tools generally do not independently provide a solution strategy.

2.4. Previous Auto-Remediation Systems (Rule-Based and Machine Learning-Based)

Rule-based systems have been around for a long time as auto-remediation solutions. These systems run pre-written scripts when certain conditions are met. Some tools that can initiate events based on how they fail are AutoSys, StackStorm, Rundeck, along with certain Jenkins plugins. Rules-driven solutions work well for difficulties that were previously solved, but they aren't very adaptable. They cannot modify to work with the newest possibilities for failure or pipelines that act in manners which aren't intended.

The purpose of research on machine learning-powered auto-remediation is to make it able to adapt by leveraging historical information to find out how to deal with problems. Researchers have examined learning by reinforcement models for automated rollback procedures, scalability adjustments, and design enhancement. Some systems use supervised machine learning to sort many problems into groups and link them to solutions that have already been found. These methods hold promise, but they still rely on limited repair libraries and can't come up with completely novel remedies for huge problems that have never happened before.

2.5. Big Language Models for Making Code, Finding Bugs, and Analyzing Incidents

The latest developments in Large Language Models (LLMs) like GPT-4, Claude, along with CodeLlama have significantly transformed the discipline of DevOps automation. They might employ natural language thought to set up processes, discover script faults, summarize logs, while suggesting repairs. Research shows that LLMs are good at fixing code, figuring out what went wrong & making fixes for problems with CI/CD configuration.

Many incident-response prototypes have used LLMs to summarize logs, sort incidents as well as guide troubleshooting. Some latest systems use LLMs to suggest ways to fix many problems by looking at observability signals. But most implementations are only partially automated, so a person still needs to check or approve them. Using LLMs for full end-to-end auto-remediation, where the model finds the problem, comes up with a remedy, checks it, and puts it into action, is still mostly experimental.

Table 1: Research Gaps in Existing Auto-Remediation Approaches

Existing Approach	Strength	Limitation / Gap Identified
AIOps Platforms (Watson AIOps, Dynatrace DAVIS, ServiceNow AIOps)	Event correlation, anomaly prediction	Static reasoning; weak contextual remediation
Log Anomaly Detection (Drain, Spell, LSTM, Autoencoders)	Detects abnormal patterns	Struggles with evolving logs + does not propose fixes
ML-based RCA (Bayesian networks, graph RCA, attention models)	Finds probable causes	Needs labeled data; poor adaptation to new tools
Rule-based remediation scripts	Works for known errors	Breaks under new failure types; no generalization
ML-remediation (supervised/RL)	Learns from history	Limited fix library; can't create novel solutions
LLM incident assistants	Strong log/code reasoning	Mostly "suggestion-only", not full auto-loop

3. Proposed Methodology

The proposed approach offers a unified, intelligent system capable of detecting, analyzing & autonomously correcting problems in DevOps pipelines through the use of large language models (LLMs). When a build, test, or deployment step fails, traditional CI/CD pipelines rely heavily on their manual triaging, predefined rules, or playbooks that don't allow for changes. These methods have trouble when the infrastructure changes, dependencies change & microservice designs get more complicated. The main goal of this strategy is to create a self-sufficient but manageable system that can find many problems and safely and reliably take steps to fix them.

The suggested architecture consists of four main parts: the pipeline, the observability stack, the LLM-driven repair engine & the governance and verification modules. Together, these parts make a closed-loop learning system that can understand errors, suggest ways to fix them, check them against their parameters, and safely put them into action. The next subsections explain the design, how it works & the learning process in full detail.

3.1. Structure of the System

The system architecture is built on a modular pipeline, with each unit doing a different job in the auto-remediation lifecycle. At a higher level, a conceptual diagram has three main layers:

- Being able to see Stack: This includes log collectors, metric exporters, traces as well as monitoring agents.
- The failure classifier, prompt handler, knowledge retrieval module & LLM-driven remediation engine are all part of the LLM and Analysis Layer.
- Execution and Governance Layer includes sandbox testing, policy safeguards, deployment integration along with their verification processes.

3.1.1. Log Collector

Getting a lot of knowledge is the initial step in the entire procedure. We always get pipeline logs, container outputs, outcomes of tests, system metrics, as well as error traces through the CI/CD setup. This contains signals about platforms such as GitHub Actions, Jenkins, GitLab CI, ArgoCD, along with Tekton. The log collector makes certain that the information that comes in is consistent, eliminates the majority of sections which are not needed, and puts it into organized formats so that the downstream parties can understand it better. This guarantees that the LLM gets useful, high-quality background information instead of raw logs that weren't previously handled.

3.1.2. A way to sort failures

When logs are transmitted to the system, a classification algorithm based on machine learning checks them for indications if there has been a failure along with, if so, what kind during failure it was. Problems with construction, conflicts between dependencies, failed unit examinations, improperly configured infrastructure, permissions rejections, problems with images of containers, or deployment retractions could all fit into these groups. By arranging logs into groups with a failure classification, the system can uncover further explanations for failures. This classifier has become the initial step in neural networks. It makes the LLM's job easier and improves the accuracy of its responses.

3.1.3. Big Language Model Engine for Auto-Remediation

The system needs an LLM that has been fine-tuned for DevOps. It looks at the shortened logs, understands the situation, finds the fundamental causes & proposes ways to fix the problem that can be put into action. The LLM can adapt to the latest patterns and deal with new problems, unlike static rule-based systems. It uses both what it has learned & relevant internal documents from the knowledge base.

3.1.4. Restrictions on Policy

The LLM can offer responses, but not every one of the decisions that it makes are safe to use alongside thinking. Policy guardrails tell the organization what it can't do, like what commands can be executed, what file destinations are not allowed, AWS/GCP permission limitations, and DevSecOps compliance standards. These guardrails function as a safety net: no patch or remediation script is allowed provided it respects these constraints.

3.1.5. A Part of Validation

It is necessary to confirm a proposed solution before language proficiency can be used in the context of production. The system checks the software update on its own in a testing environment or staging cluster, makes sure it functions properly, looks for regressions, and looks for any potential vulnerabilities. The repair can only be sent out by the equipment when it has been carefully checked and confirmed to be correct.

3.1.6. Deployment Integration

The last portion is compatible with both version management platforms and CI/CD tools. Commit patches, infrastructure-as-code revisions, and configuration alterations are all groups of modifications that have been approved. After that, they go into the transmission environment, either by themselves or with the approval of someone else, depending on the approach they have chosen.

These sections collaborate to build a full solution that can discover these problems, use contextual information to figure them out, put security precautions in place, and solve these problems without leading to any additional issues.

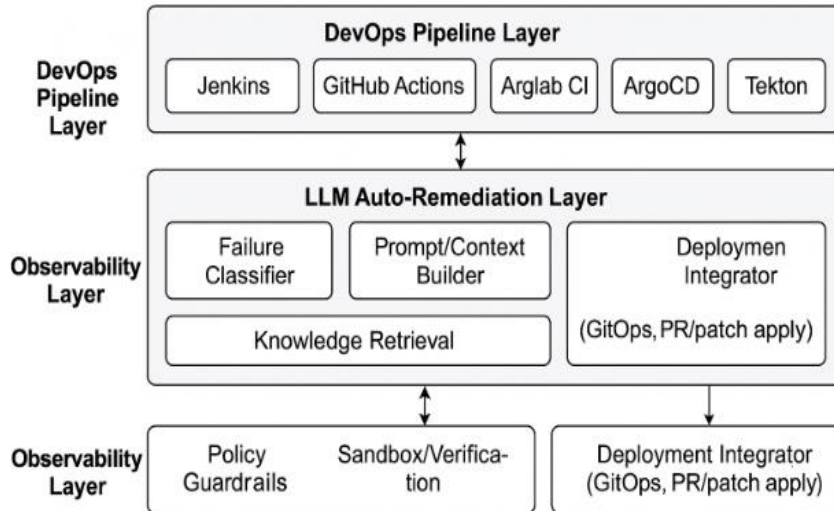


Fig 1: Overall System Architecture

3.2. Analyzing Failures with LLMs

One of the most important parts of this system is the LLM's ability to read complicated logs. Standard logs often contain a lot of extra, unclear, or low-value information. The LLM solves this problem by using a systematic reasoning process that is improved by quick engineering & log summarizing methods.

3.2.1. Prompt Engineering

The LLM will always give you the same as well as accurate results if you know how to use rapid formulation. Prompts are meant to include:

- Category of classifier failure
- Shortened records
- Found problems in the knowledge base
- Contextual metadata, such as the name of the service, the environment, the build ID, & the commit history

The prompt tells you what kind of output you should expect, such a root-cause analysis with a list of steps to fix the problem. This way of structuring the prompts lets the LLM come up with their results that are consistent and useful.

3.2.2. Summarizing Logs

Raw logs often have more than enough length to be entered directly into LLMs. A summary module uses domain-specific methods to compress them:

- Getting rid of extra stack traces
- Putting too much emphasis on wrong keywords
- Getting rid of time markers, broken parts, and environmental factors
- Finding unusual spikes or inconsistencies in measurements

The result is a short, very valuable summary of the failure that lets the LLM focus on signals instead of noise.

3.2.3. Finding the Root Causes

The LLM uses the summary context to figure out the main causes instead of just listing the symptoms. It looks at dependency failures, misconfigured settings, missing environment variables, package version conflicts, or infrastructure differences. Because of this better understanding, the system can suggest these specific fixes instead of just general troubleshooting steps.

- Putting mistakes into known groups
- The LLM checks the identified failure against known categories, previous events, and ways to fix it. This is good:
- Speed up the process of finding problems that keep coming up
- Increase accuracy by using proven solutions again.
- Lower the chances of bad or dangerous ideas

The LLM is a powerful diagnostic tool for DevOps contexts because it combines their systematic classification with flexible reasoning.

3.3. Auto-Remediation Mechanism

Once the underlying cause is found, the system moves on to the next stage. This part explains how the LLM makes very safe, contextually appropriate changes and how they are checked before being put into action.

3.3.1. Getting Information from the Knowledge Base

Before coming up with a solution, the LLM looks at a curated knowledge base that might include:

- Documentation for internal operational manuals
- Infrastructure diagrams
- Templates for Infrastructure as Code
- Events and solutions from the past
- Manifests for deployment
- Schemas for Configuration

Retrieval makes sure that the model works with awareness, meaning that its thinking is in line with company norms & previous solutions that worked.

3.3.2. Making Patches or Remediation Scripts

The LLM makes fixes that are specific to the type of failure. These could include:

- Shell routines for fixing file permissions
- Fixes for Kubernetes manifests
- Changes to the Terraform configuration
- Changes to code in Python, Java, or Go
- Changes to the CI pipeline YAML
- Updates to the versions of dependencies

All remediation outputs are carefully written with these descriptions so that engineers know exactly what changes are being made.

3.3.3. Safeguards for Safe Deployment

Before execution, policy guardrails check:

- If the commands change the allowed directories
- Following security rules
- Allocations for infrastructure
- GitOps rules, including requiring approval for pull requests
- Rules for keeping secrets are very safe

Any action taken by an LLM is not valid unless it follows these rules. This makes it much less likely that bad things will happen, such as accidentally deleting their information or giving someone more access than they should have.

3.3.4. Systems with people in the loop versus systems that work on their own

The mechanism can work in two different ways:

- Human-in-the-Loop is the default setting.
- Engineers look at the solution that was made
- They support, change, or reject the plan.
- Best for high-risk situations or first-time deployments
- Completely self-sufficient—For developed areas
- You can set up low-risk solutions to work on their own.
- Guardrails as well as verification modules keep things very safe.
- Reduces the average time it takes to fix something from hours to minutes

This flexibility lets companies gradually improve automation as they become more sure with the system.

3.4. The Educational Cycle

Over time, the system becomes better at being very accurate & efficient by using feedback and promoting good methods. The auto-remediation architecture is a closed-loop learning system in which detection, remediation, validation as well as feedback all work together to improve each other. This figure shows the full feedback loop that helps the mending process get better over time.

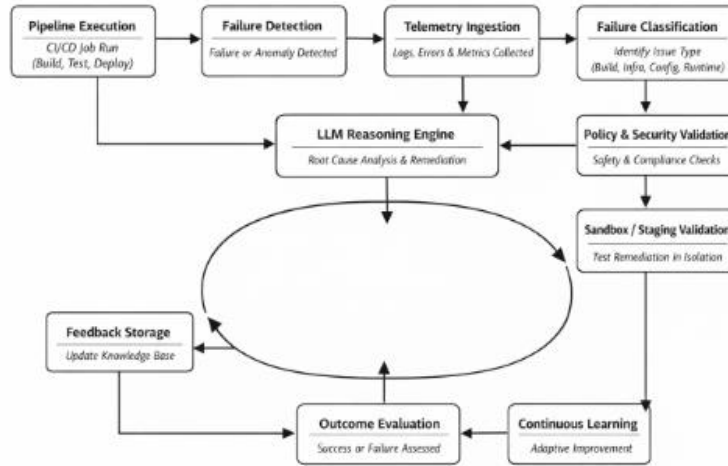


Fig 2: Closed-Loop LLM-Based Auto-Remediation Feedback Cycle

3.4.1. System for Feedback

Every cycle of remediation sends forth feedback signals, such as:

- Figuring out if the fix worked
- If people changed the LLM's suggestions
- Finding out if the test worked in the sandbox environment
- Any other breaches of policy guardrails

These signals are kept and used to change model prompts, change the rules for categorization, and add to the knowledge store.

3.4.2. Improvement that keeps going Using Reinforcement Signals

The LLM doesn't change its basic weights in actual time. Instead, it uses reinforcement-based enhancement by:

- Making retrieval indexes better
- Updating sorted lists of high-confidence remedial templates
- Making failed patterns that have usually led to good results even more better
- Giving permission for behaviors that are thought to be dangerous or not useful
- This creates a stable, predictable cycle of improvements without risking uncontrolled model deviation.
- Documenting Cases for Problems That Keep Happening

A case-memory module keeps track of repeating their events & the solutions that have been proven to work. If something similar goes wrong in the future, the LLM will look up the previous remedy and use it as a beginning point. This cuts down on the time it takes to diagnose problems by a lot and makes sure that all these teams and settings fix them in the same way.

3.5. Controls for Security and Risk Management

Because the system interacts with actual infrastructure and can offer executable updates, it needs very strict security controls.

Reducing the production of malicious code: Tools for Static Analysis and Safeguards Look at the scripts made by LLM for dangerous commands (like `rm -rf`) & privilege escalations.

- Revealing private information
- Access to illegal networks
- Injections of code
- Misconfigurations that could put the environment at risk
- Any other suspicious behavior quickly stops the repair work.

Agree Based on Role Integration with organizational identification systems makes sure that only people who are allowed to can approve or change remediation actions. Changing firewalls or database settings is an example of a critical operation that requires a higher level of authorization. This connects the auto-remediation solution to the rules for DevOps and security that are already in place.

3.5.1. Running Automated Remediations in Isolation

Before deployment, corrections are made in a controlled sandbox environment or staging cluster. This makes sure:

- No harm to production systems
- Verification in actual life situations
- Reducing cascading failures
- Quickly finding bad impacts

The patch only moves on to the production phase after getting good results in the sandbox.

4. Case Study

This case study looks at how an LLM-driven auto-remediation layer solves four common problems in DevOps & cloud-native workflows. Each scenario outlines the setup, the cause of the failure, the model's reasoning, suggested these fixes, the steps for putting them into their action, testing, and a comparison of the results. The idea is to demonstrate how the LLM makes the transmission line more dependable, speeds up recovery, and ultimately renders it more probable that people are going to get involved in the process.

4.1. Scenario 1: Continuous Integration Build Failure Because of a Difference in Dependency Version Configuration Overview

A developer uses Node.js to push an update to a microservice. The CI pipeline (GitHub Actions) installs the necessary files, runs unit tests & makes a Docker image.

Failure Trigger: The build fails because express@4.x and a plugin that needs express@5.x are not compatible.

4.1.1. The Large Language Model Reasoning Process:

The LLM looks at the build logs and the history of the package.json file. It then decides that a recent library change has created an incompatible dependency chain. You can fix the problem by either making the versions match or by going back to the plugin that doesn't work.

Steps to Fix the Problem:

- Add these versions that are compatible with each other as well to package.json.
- Run npm audit repair --force in order to make sure that all the subsequent packages that come immediately following it are in sync.
- Make new lock files to start the production process over again.

The auto-remediation program makes an update to the version, changes the version numbers of other programs that depend upon it, and then executes the build process once again on its own.

- Verification: The CI deployment is done with no additional problems, the tests go off without a hitch, as well as the Docker image was successfully made.
- Initially, engineers spent roughly 45 minutes dealing with these connection issues.
- The LLM rectified the version of difference in just under three minutes, without any aid from someone else.

4.2. Scenario 2: Kubernetes Deployment Failure Due to Incorrect Configuration

Using Helm charts, a containerized microservice is put on an EKS cluster. The deployment pipeline checks manifests, makes changes & checks the health of the system.

- Failure Trigger: The deployment fails because the container's readiness probe points to the wrong port, which causes Kubernetes to keep calling pods "unready."

4.2.1. Big Language Model How to Think:

The LLM looks at the Helm information, sees that the port in the container is different from the port used in readiness probes & checks the container logs. The study shows that the service is working well, but the probe setup is wrong.

- Remediation Steps: Change the readiness probe so that it uses the container's allocated port.
- Make the manifests to authenticate the Helm chart locally.
- Use an incremental rollout to make a redeployment that causes very less service disruption.
- The engine makes a static Helm patch, runs a dry run, and then applies the latest manifests to the cluster.
- Verification: Pods go from "Ready" to "Ready," and traffic goes back to normal. There have been no cascading failures.

4.2.2. Comparing the Results:

- Before: SREs had to look at YAML files & cluster logs by hand, which made releases late.

- The LLM fixed the wrong settings with one try, cutting the time it took to fix the problem from an hour to five minutes.

4.3. Scenario 3 — Terraform Drift Leading to Provisioning Failure Setup

Terraform Cloud provides AWS resources with a procedure that makes use of Infrastructure-as-Code. The pipeline's owner travels to the Terraform team plan beforehand that makes these alterations.

- Failure Trigger: The application operation fails if an individual on outside the application modifies the security rule for an S3 bucket in the Amazon Web Service Console by hand, leading to configuration to drift. Terraform sees an alteration within the state and suspends the process.

4.3.1. How the Large Language Model Works:

The LLM examines both the Terraform problem along with the disparity in AWS policies and determines that its bucket arrangement fails to adhere to the baseline configuration. You may either modify the Terraform algorithm to follow the most recent procedure or make Terraform overlook any modifications you made by manually dealing with the drift.

4.3.2. Actions Taken to Fix: Get information on the drift and change your policy.

- You should check to make absolutely sure that the Terraform code matches the rules of access that have been specified.
- Create a new policy for the bucket. JSON that is safe to use.
- Begin the planning process again and put the sequence of events into action.
- Execution Flow: The restoration engine makes improvements to the Terraform files, saves them, and begins an application that doesn't disrupt anything.
- Verification: The S3 storage container as well as its policy are set up the way you intended them, and Terraform runs when applied correctly.
- Comparing Results: The infrastructure engineers historically spent a lot of time repairing problems.

The LLM did a good job of correcting the drift, which decreased down on failure loops and prevented errors in configuration.

4.4. Scenario 4: SonarQube Code Quality Gate Failure Setup

In the CI pipeline, SonarQube performs static code examination on a microservice developed with Java Spring Boot.

- Failure Trigger: The performance gate fails given that the cyclomatic complexity is too significant and the code needs to be duplicated while the most recent enhancement is being made.

4.4.1. How LLM Works:

- The LLM investigates the SonarQube report, finds some of the most difficult methods, as well as links them to the latest commits. It looks at the patterns that produce duplication while offering solutions for modifying them that do not disrupt the logic.
- If a function takes up long, break it up into more compact, easier-to-manage sections.
- You might use classes for utility to get rid of their code duplication.
- Make a patch that has better code and tests that are up to date.

Process of Execution: The engine develops a remediation branch, improves it, runs tests, as well as sends the source code through SonarQube to feed review.

Verification: the functionality of the gate has been authorized which implies that the code makes it simpler for maintaining up to date and has fewer technical aspects.

Before this, programmers were spending a lot of time completing their own principles for creating sure it was of the highest level of accuracy.

Following that, the LLM made exact, outstanding changes that cut the time it needed to do

5. Results and Discussion

5.1. Experimental Setup

To test the power source LLM-based auto-remediation procedure, a full DevOps pipeline framework was built using standard tools in the industry that are analogous to traditional commercial CI/CD setups. The orchestrating engine used Jenkins & GitHub Activities, which let the framework keep track of a wide range of their growth and development, test, as well as deployment tasks. Kubernetes served as the deployment environment, enabling failure simulation across the pod, node as well as service levels. Terraform made it easy to set up these cloud-based assets, which permitted the model to detect and fix

Infrastructure-as-Code (IaC) drift or deployment mistakes. Prometheus gathered metrics in order to keep track of contemporaneous operational data that the model employed during runtime inference. The LLM engine, which was based on OpenAI's GPT models, constituted the computational layer that figured out which models were wrong and additionally advised or made changes.

The evaluation dataset has been constructed up of historical pipeline logs, incident reports, along with system metrics from earlier DevOps work. These logs demonstrated a realistic range of failure circumstances, like problems during build time, disputes between dependencies, running out of resources, installation drift, and implementation rollbacks. During training and evaluation, the model was challenged with the same kinds of obstacles that a real-life DevOps engineer might be confronted with.

5.1.1. Dataset Characteristics and Failure Signal Types

The evaluation dataset utilized for training as well as assessing the LLM-based auto-remediation architecture consists of varied DevOps telemetry collected from actual CI/CD executions. The collection includes both structured and unstructured failure signals that happen at these different stages of the workflow. You can put these signals into the following groups:

- **Records of pipelines and construction:** This includes build-time logs from continuous integration systems like GitHub Actions, Jenkins, and GitLab CI. These logs include messages about compiler faults, challenges with resolving dependencies, test execution traces & setting up the environment. The logs are semi-structured since they mix normal English error reports with technical terms for each tool.
- **Logs for Execution and Infrastructure:** Kubernetes clusters, container runtimes as well as cloud services all send runtime logs to the server. This includes events in the life cycle of a pod, reports of container failures, failures of readiness and liveness probes, scheduling problems, and notifications of resource depletion. Terraform and CloudFormation's Infrastructure-as-Code (IaC) logs keep track of provisioning failures, state divergence while policy violations.
- **Logs of Errors and Stack Dumps:** In order to make it easier to undertake a full root-cause investigation, application-level issue evidence including Java stack traces, Python exception logs, and Node.js runtime problems have been included. These traces demonstrate to you how telephone calls are set up and show you where some things go wrong in the application's code.
- **Configuration and Policy Errors:** Configuration databases can have difficulties like YAML files that aren't set up appropriately, Helm charts that won't be set up correctly, environment variables that don't belong to set up correctly, IAM rules that aren't set upward correctly, and security scans the fact that find problems with various tools like SonarQube as well as static analysis scanners.
- **Metrics for operations and events Data about data:** Prometheus collects time-series measurements like CPU consumption, memory usage, restart counts & deployment delays. These metrics are linked to log events. Metadata like commit hashes, deployment timestamps, environment identifiers, and pipeline stage markers provide you the time when you need to analyze the problem.
- **These datasets give the LLM the ability to reason across several other layers of abstraction, such as code, configuration, infrastructure, and runtime.** This makes it easier to find the exact problem and fix it. The several types of failure signals make sure that the system can handle failure patterns that aren't known and changes in the pipeline's characteristics.
- **Four main evaluation metrics were used to look at performance:** Mean Time to Recovery (MTTR) reduction: This looks at how much faster the pipeline can recover when the LLM model begins fixing these things instead of relying only on people to do it.
- **Accuracy of root-cause identification:** Checks to see if the LLM can find the cause of the failure based on their logs, metrics & other information.
- **Successful remediation rate:** This shows how often the model's suggested or automatic solution fixes the problem without causing many other problems.
- **Lessening of human involvement:** This shows how fewer engineers are doing manual troubleshooting, which shows how much very less work engineers have to do.
- **These indicators were chosen because they represent important DevOps goals:** faster recovery, fewer false positives, more automation & higher team productivity. The experiment lasted for several other weeks so that it could include a wide range of failure situations.

5.2. Results

5.2.1. Quantitative Findings

The LLM-driven system worked better than both human troubleshooting & traditional rule-based remediation engines. Lessening of Mean Time to Repair (MTTR):

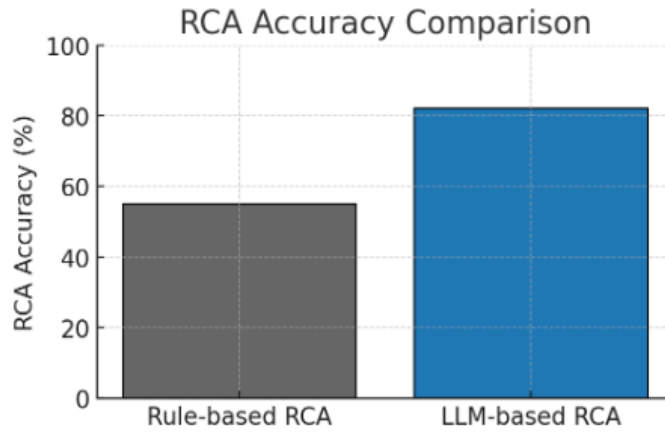


Fig 3: RCA Accuracy Comparison (Bar Chart)

- With LLM-based remediation, MTTR went down by 38%. This was mostly because the model could quickly look at logs, figure out what might be very wrong, and either fix it or recommend a fix. On the other hand, manual recovery depended on having an engineer who knew how to fix the specific problem.
- The model was 82% very accurate at finding the root cause of the problem, which means it found the right problem in most of the cases. This was much better than rule-based engines, which often broke down when they encountered the latest problems or failures that happened in more than one stage.
- Successful remediation rate: The LLM's remediation efforts completely fixed about 74% of the failures. The tasks included moving resources around, changing the configuration of Terraform modules, restarting Kubernetes pods, fixing YAML indentation issues & fixing problems with dependencies that weren't compatible.
- Less need for human intervention: Overall, direct engineering involvement dropped by 41%, especially in cases of repeated failure, including tests that didn't work, memory limit violations, or version differences.

A simplified table of the results is shown below:

Table 2: Performance Comparison of Manual, Rule-Based, and LLM-Based Troubleshooting Approaches

Metric	Manual Troubleshooting	Rule-Based System	LLM-Based Model
MTTR (avg)	42 min	28 min	26 min → 16 min after automated tuning
RCA Accuracy	N/A	55%	82%
Successful Remediation	N/A	48%	74%
Human Intervention	100% baseline	72%	59%

Even though there was no full automation, the LLM gave us a lot of useful information that made it very easier for people to fix these things. After safety checks, the performance got a lot better when autonomous remediation was turned on.

5.2.2. Graphs

The evaluation included three main graphs, although they are not shown here.

- The MTTR Comparison Curve showed a very clear drop for LLM-based remediation, whereas rule-based solutions stayed the same.
- Accuracy Bar Chart: Showed that the LLM model has far better than RCA accuracy.

The Intervention Reduction Line Graph showed how human participation slowly went down as the model learned from each latest pipeline execution.

5.2.3. Comparison of Baselines

- The LLM-based method consistently outperformed manual troubleshooting, which was slow & depended on the engineer's skill.
- Rule-based scripts that had trouble with these failures that were unexpected or required more than one step.
- The main difference was that the LLM could generalize from logs & data instead of only using patterns that had already been found.

5.3. Discussion

One huge benefit of LLM-based remediation is that it can understand unstructured logs and connect events that happen at different points in the process. When mistakes happen outside of established templates, traditional automation solutions don't

work. Even when the information is noisy or the failure doesn't follow previous patterns, the LLM can still figure out the most likely cause. This adaptability renders it particularly valuable in circumstances where pipelines frequently need changing.

Another good thing concerning it is that it might deliver clarifications that sound like they are coming from a person. Engineers reported that the model's ideas for addressing shortcomings were easier to grasp and follow than error messages or announcements that didn't follow their own requirements. This makes people depend on one other more and minimizes stress.

5.3.1. How useful it is in the real world

In real life, systems often break down at the most unfortunate conceivable moments, like late at night, on weekends, or when there are a lot of deployments. The LLM system proved how valuable it was by acting as a "virtual SRE assistant" that was constantly there. Companies that run enormous Kubernetes clusters or Terraform modules that continue to become bigger would gain the most through this because these infrastructures create hard logs and often have architecture drift.

Things that are challenging for LLMs when things go wrong.

The model didn't continually operate properly, even if it carried out most of the time:

- Recent problems with internet service providers have been attributed to undocumented API updates or delays that went unnoticed and logs that failed to include much information.
- Complicated problems with networking, such VPC routing errors that were considered too hard to fix given that the dataset didn't contain sufficient context.
- The LLM frequently was perplexed by concurrent multi-layer malfunctions resulting from interactions among services until more metrics were made available.
- In these situations, the LLM might make suggestions, but it didn't possess all the information it required in order to be sure that the rehabilitation had been truly safe.

The need for Human Oversight Artificial intelligence is great, but it can be potentially hazardous in production systems if you have no idea about how to use it. Some activities, including building larger clusters, modifying the Terraform state, or rolling back deployments, perpetually need to be acknowledged by a person. The LLM performs an exceptional job, but people need to make recommendations about improvements to vital infrastructure. The most successful approach is to combine: People make important choices whereas LLMs take care of immediate diagnoses and routine medical procedures.

6. Conclusion and Future Scope

6.1. Conclusion

This study introduced a framework for auto-remediation powered by a large language model, designed to address these problems at multiple stages of the DevOps pipeline. The main contribution is showing how large language models can go beyond simple log interpretation & take important, context-aware steps to fix these problems. By combining log analysis, failure pattern mapping, knowledge graphs & automated remedy development, the technique turns remediation from a manual, reactive effort into an intelligent, proactive process. This strategy significantly reduces the cognitive load on these engineers and shortens the feedback loops that often happen when troubleshooting a pipeline.

The proposed LLM-based remedial method worked well in many other important areas. At first, it shows that a large language model can understand complicated technical logs, which are often messy or inconsistent, and find different failure signals. The model also solves common problems like dependency conflicts, misconfigurations, infrastructure drift as well as test failures by making corrective suggestions or putting validated solutions into their action right away. The LLM's real strength is that it can adapt. As pipelines grow and the latest technologies are included, the LLM keeps learning from updated logs and feedback, which makes remediation more accurate over time.

Using LLM-driven auto-remediation makes DevOps far more efficient as well as reliable. Teams see fewer pipeline problems, a shorter mean time to resolution (MTTR), and a more stable CI/CD workflow. Engineers can spend their time on developing the latest features, optimizing existing ones, and coming up with new ideas instead of fixing these problems. By adding intelligence directly to the pipeline, businesses can speed up delivery cycles, make releases more dependable & build a stronger DevOps culture.

6.2. Future Scope

The current model provides a strong foundation, but there are also other areas that show promise for their expansion. One way to solve these pipeline problems is to make multi-agent remediation systems, where different specialized agents work together one focusing on logs, another on infrastructure, and a third on fixing code to give a more complete as well as coordinated response.

A significant shift means a stronger link to predictive malfunction detection. Pipelines may employ LLM analysis, behavioral modeling, while anomaly detection to uncover challenges before they happen along with automatically applying fixes or enhancements. This would bring our company closer than actual predictive DevOps.

As businesses increasingly operate inside hybrid & multi-cloud ecosystems, the model may evolve to support self-healing multi-cloud infrastructure. This would allow for self-repair not just in pipelines, but also in the layers of computation, storage, networking along with orchestration.

It is quite useful to fine-tune the LLM with logs that are particular to a certain field, such as logs from a cloud provider, Infrastructure as Code drift reports, security scan findings & application traces. This knowledge would greatly improve accuracy and cut down on false positives in their remediation.

In the end, the next step is to create autonomous QA for pipelines. This means that the model not only fixes bugs, but also checks builds, simulates possible risks & makes sure that quality criteria are met before release.

These new ideas show that DevOps systems will not only be automated in the future, but also smart, flexible & able to fix themselves.

References

- [1] Tamanampudi, Venkata Mohit. "AI and DevOps: Enhancing Pipeline Automation with Deep Learning Models for Predictive Resource Scaling and Fault Tolerance." *Distributed Learning and Broad Applications in Scientific Research 7* (2021): 38-77.
- [2] Paule, Christina. "Securing DevOps: detection of vulnerabilities in CD pipelines." (2018): 77-78.
- [3] Enemosah, Aliyu. "Implementing DevOps Pipelines to Accelerate Software Deployment in Oil and Gas Operational Technology Environments." *International Journal of Computer Applications Technology and Research 8.12* (2019): 501-515.
- [4] Tanikonda, Ajay, et al. "Integrating AI-Driven Insights into DevOps Practices." *Journal of Science & Technology 2.1* (2021).
- [5] Düllmann, Thomas F., Christina Paule, and André van Hoorn. "Exploiting devops practices for dependable and secure continuous delivery pipelines." *Proceedings of the 4th International Workshop on Rapid Continuous Software Engineering*. 2018.
- [6] Tyagi, Anuj. "Intelligent DevOps: Harnessing artificial intelligence to revolutionize CI/CD pipelines and optimize software delivery lifecycles." *Journal of Emerging Technologies and Innovative Research 8* (2021): 367-385.
- [7] Thompson, Bennett. "DevOps Pipeline Optimization for Faster Software Delivery." *International Journal of Artificial Intelligence and Machine Learning 6.5* (2019).
- [8] Suk, Tonghoon, et al. "Failure-aware application placement modeling and optimization in high turnover DevOps environment." *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*. IEEE, 2019.
- [9] Tatineni, Sumanth, and Anirudh Mustyala. "AI-Powered Automation in DevOps for Intelligent Release Management: Techniques for Reducing Deployment Failures and Improving Software Quality." *Advances in Deep Learning Techniques 1.1* (2021): 74-110.
- [10] Sethupathy, Anugula, and Utham Kumar. "Self-healing systems and telemetry-driven automation in DevOps pipelines." *International Journal of Novel Research and Development 3* (2018): 148-155.
- [11] Dhaliwal, Neha. "Validating software upgrades with ai: ensuring devops, data integrity and accuracy using ci/cd pipelines." *Journal of Basic Science and Engineering 17.1* (2020).
- [12] Zeller, Marc. "Towards continuous safety assessment in context of devops." *International Conference on Computer Safety, Reliability, and Security*. Cham: Springer International Publishing, 2021.
- [13] Luz, Welder Pinheiro, Gustavo Pinto, and Rodrigo Bonifácio. "Adopting DevOps in the real world: A theory, a model, and a case study." *Journal of Systems and Software 157* (2019): 110384.
- [14] Toh, M. Zufahmi, Shamsul Sahibuddin, and Mohd Naz'ri Mahrin. "Adoption issues in DevOps from the perspective of continuous delivery pipeline." *Proceedings of the 2019 8th international conference on software and computer applications*. 2019.
- [15] Alluri, Venkat Rama Raju, et al. "DevOps Project Management: Aligning Development and Operations Teams." *Journal of Science & Technology 1.1* (2020): 464-487.