

AI-Enabled DevOps and Automated CI/CD Pipelines for Reliable Cloud Service Delivery in Saudi Arabia

Mohammed Vaquar Ahmed Haseeb^{1*} 

¹Independent Researcher

*Corresponding Author

Mohammed Vaquar Ahmed Haseeb

Independent Researcher

Article History

Received: 19.08.2026

Accepted: 12.10.2026

Published: 14.10.2026

Abstract: The reliability of cloud services is now increasingly determined by how safely organizations can modify their software, rather than just by the redundancy of their infrastructure. Although continuous integration and continuous delivery pipelines have shortened the time between releases, problems such as static rules, fragmented telemetry, manual approval hurdles, and a lack of feedback between deployment and operations can cause failures to become worse as cloud systems become more distributed. This study looks at the way in which artificial intelligence can be incorporated into DevOps and into automated CI/CD pipelines in order to enhance reliable cloud service delivery in Saudi Arabia. A structured integrative review of peer-reviewed research published between 2020 and 2025 was carried out in the fields of software engineering, AI-assisted testing, AIOps, observability, cloud automation, and literature on Saudi cloud adoption. Thirty studies, verified by their DOIs, were analysed using five analytical dimensions: pipeline stage, AI function, reliability mechanism, governance requirement, and relevance to Saudi enterprise adoption. The evidence shows that AI provides the most significant improvement in reliability when it functions as a bounded decision-making layer based on high-quality telemetry, rather than as an unbounded substitute for engineering judgment. Together, predictive detection of build failures, intelligent regression testing, deployment-risk scoring, anomaly detection, root-cause localisation, and feedback-driven rollback can collectively reduce the risk resulting from frequent releases. Nevertheless, the benefits in terms of reliability depend on the use of traceable data, policy controls, progressive delivery, explainable decisions, and human authority for important changes. The review puts forward a closed-loop reference model that connects code changes, automated verification, risk assessment, controlled release, runtime observability, and learning. For Saudi enterprises, the model stresses the importance of trust, service-level assurance, governance, organizational readiness, and the adoption of locally appropriate cloud operating practices. The paper ends with specific implementation priorities and a research programme aimed at validating the approach on a production scale.

Keywords: Artificial Intelligence; DevOps; Continuous Integration and Continuous Delivery (CI/CD); Cloud Service Reliability; Automated Software Testing; AIOps; Predictive Failure Detection.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC BY-NC 4.0) which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

Citation: Haseeb, M. V. A. (2026). AI-Enabled DevOps and Automated CI/CD Pipelines for Reliable Cloud Service Delivery in Saudi Arabia; *Glob Acad J Econ Buss*, 8(5), 1875-1886. <https://doi.org/10.36348/gajeb.2026.v08i05.072>

1. INTRODUCTION

The way in which software is delivered in a cloud environment has altered the unit of work in reliability engineering. In the case of traditional release methods, teams could divide development, testing, deployment, and operations into fairly separate stages. In contrast, modern DevOps combines these activities into a continuous process where each commit can set off automated actions relating to building, testing, security, packaging, and deployment. Artificial intelligence is being used more and more within this process since software engineering already produces large amounts of code, logs, traces, test results, configuration histories, and records of incidents that can be used for prediction and optimisation [1]. Research into software engineering as applied to AI-based systems also indicates that dependability, testing, and quality assurance are key issues when intelligent components are included in operational software [2]. The opportunity that arises is not merely to automate additional steps, but to make the automation sufficiently context-aware so that it can tell the difference between ordinary changes and those which pose a threat to service reliability.

The body of evidence relating to the use of AI in software engineering has grown very quickly. Reviews of deep learning and machine learning indicate that these techniques are being widely applied to areas such as defect prediction, testing, maintenance, and other activities throughout the software lifecycle, yet they also point out ongoing weaknesses in terms of reproducibility, industrial validation, and data quality [3, 4]. The significance of these weaknesses lies in the fact that in CI/CD an incorrect model recommendation can lead to an operational action being taken within a matter of minutes. If the system predicts that a test can be omitted, that a deployment is low risk, or that an anomaly is harmless, then the likelihood of a service failure is altered. AI-enabled DevOps should therefore be viewed as a control system not as a set of separate productivity tools. The quality of such a system depends on the data used in the model, the limits placed on automated actions, the ability to observe the results, and the capacity to halt or reverse a release when new evidence emerges.

The problem of having to exercise control becomes more severe with cloud-native architecture since microservices, containers, APIs, infrastructure-as-code, managed cloud services, and workloads that scale dynamically all increase the number of components and dependencies which can change independently. Research into intelligent cloud-native architecture connects DevOps practices with automated delivery, elastic infrastructure, and operational feedback [5]. Empirical studies on CI/CD

also show how continuous practices can reduce the delays in coordination between development and operations, although the benefits depend on integrated processes and toolchains rather than simply having a pipeline [6]. GitOps applies the same idea by considering the desired operational state as a version-controlled, declarative configuration, thus improving traceability and reproducibility [7]. Although these methods provide a solid foundation for automation, they do not in themselves ensure that a particular change is safe given the current runtime conditions.

This point is particularly significant for Saudi companies, since alongside the growth of digital services their adoption of cloud computing is progressing, together with the need for data governance and their growing reliance on customer and internal systems that are always available. Failures in reliability can have an impact on public services, financial processes, industrial operations, retail channels, and corporate productivity. Given this, the situation in Saudi Arabia calls for an examination that brings together three areas which are usually handled separately: the use of AI in software engineering, DevOps automation, and the assurance of cloud services. The key issue is not whether AI can produce code or improve a step in a pipeline; rather it is how decisions supported by AI can be incorporated throughout the delivery lifecycle in such a way that release speed is increased without passing on unacceptable operational risk to production.

The review describes reliable delivery as forming a closed loop in Saudi enterprises, since change evidence, automated verification, AI risk assessment, staged deployment, runtime telemetry, incident outcomes, and human overrides all continuously feed information to guide decisions.

2. Aim, Objectives and Review Questions

This review seeks to combine the latest evidence regarding the way in which AI can be incorporated into the automated DevOps and CI/CD control cycles in order to enhance the reliability of cloud-service delivery within Saudi companies. The analysis is guided by five objectives. The first of these is to identify the points at which AI can be introduced throughout the various stages of planning, building, testing, releasing, deployment, and runtime operations, depending on the reliability function that it supports. The second objective is to look at how historical pipeline data and real-time telemetry can be turned into predictive controls for managing release risk, prioritising tests, detecting anomalies, and carrying out recovery. The third involves assessing the safeguards needed to avoid automation leading to the emergence of new failure modes, such

as progressive delivery, policy gates, observability, auditability, and the possibility of human override. The fourth objective is to interpret the technical findings in light of factors related to cloud adoption in Saudi Arabia, such as trust, service assurance, organisational readiness, and governance. The fifth is to establish specific, measurable priorities for implementation and to identify the research gaps that need to be addressed in order to achieve large-scale adoption.

The review therefore poses four questions: What are the AI-supported activities that have the most clear link with reliable CI/CD delivery? How should these activities be integrated into an end-to-end reliability loop rather than being used as separate tools? What technical and governance conditions are needed for bounded automation in enterprise cloud environments? And which contextual factors should Saudi organisations take into account when applying international evidence to operational practice?

3. REVIEW METHODOLOGY

A structured integrative review was chosen since the subject area involves empirical software-engineering studies, systems research, secondary reviews, AI-assisted testing, cloud operations, and research into national adoption. The time frame for the review was limited to the period 2020–2025 in order that the evidence might reflect present-day cloud-native practices and take into account the fast development of machine learning and large language models. The search terms combined references to artificial intelligence, machine learning, large language models, generative AI, and AIOps with those for DevOps, CI/CD, continuous integration, continuous delivery, GitOps, cloud service, testing, observability, anomaly detection, deployment, reliability, and resilience. Where necessary, the searches were supplemented by carrying out backward and forward checks on the most relevant studies.

A study was included in the review if it had been peer-reviewed, was published within the specified time frame, and made a direct link between AI or automated delivery practices and software quality, deployment, monitoring, cloud operation, or reliability. Empirical studies, systematic reviews, architecture papers, and well-described frameworks were taken into account provided that they contained adequate methodological or operational details for synthesis. However, studies were excluded if they were outside of the defined period, were duplicates, were not related to a continuous-delivery or cloud-

service environment, or did not offer enough evidence to allow the proposed mechanism to be interpreted. Even if CI/CD was not the main topic of the study, Saudi-focused sources were kept in the review if they provided evidence regarding cloud trust, secure cloud-edge workflows, or AI-adoption readiness, since these factors influence the feasibility of implementation.

The set of final evidence included thirty items, and the DOI entries for these were verified against publisher or bibliographic metadata. When assessing the quality, five criteria—adapted from software-engineering review practices—were applied: the clarity of the research purpose, the appropriateness of the method, the adequacy of the operational context, the credibility of the evidence, and the explicit consideration of limitations. Instead of giving a simple rating to the various types of study design, the criteria were used to determine whether a finding was adequately supported for inclusion in a thematic synthesis. In cases where a paper put forward a conceptual capability without having validated it through production, the review regarded that capability as a research direction rather than as established operational performance.

Five analytical dimensions were used in the data extraction. The first one noted the CI/CD stage which was influenced by the intervention. The second one dealt with the AI task, for example failure prediction, test generation, anomaly detection, root-cause analysis, configuration recommendation, or resource forecasting. The third identified the reliability mechanism, this being defect prevention, a reduced blast radius, faster diagnosis, faster recovery, or improved visibility. The fourth recorded the governance requirements, particularly explainability, approval, provenance, policy enforcement, and retraining controls. The fifth evaluated the transferability to Saudi enterprises, focusing on cloud trust, service-level assurance, skills, security, and organizational maturity.

The synthesis was thematic not meta-analytic since the studies make use of different datasets, system architectures, evaluation metrics, and definitions of reliability. The analysis thus looks at the mechanisms and patterns of evidence rather than combining incompatible effect sizes. Making this choice also avoids the situation where improvements achieved in a specific testing or anomaly-detection task could be presented as direct evidence of enterprise service reliability. Instead, the review examines whether each capability can contribute to a defensible control loop in which automated decisions are observable, reversible, and accountable.

4. Evidence Synthesis

Table 1: Evidence synthesis of AI functions across the CI/CD reliability chain

Pipeline stage	AI role	Reliability contribution	Primary control
Change / planning	Change-risk classification; dependency impact analysis	Targets verification effort and identifies high-risk components early	Versioned inputs and explainable features
Build	Build-failure and resource prediction	Detects unusual pipeline behavior before release	Deterministic build and artifact integrity
Test	Flaky-test prediction; test generation; test prioritization	Improves verification efficiency and defect discovery	Execution-based validation and periodic full suites
Release	Deployment-risk scoring; evidence aggregation	Selects release strategy according to measured risk	Policy thresholds and human approval for high impact
Deploy	Canary analysis; progressive traffic decisions	Limits blast radius and supports early rollback	SLO-based promotion and reversible changes
Runtime	Anomaly detection; root-cause localization	Accelerates detection, diagnosis, and recovery	Correlated telemetry and auditable remediation

4.1 From Static Pipelines to Risk-Aware Delivery

Traditional pipelines are deterministic in that when the specified conditions are fulfilled, the next task begins. While this is beneficial for ensuring reproducibility, the static logic proves to be fragile whenever repositories, dependency graphs, infrastructure states, and workload conditions change more quickly than the pipeline rules are updated. Research in the field of AIOps shows a different approach in which telemetry is examined to identify abnormal states either before or during production. Fawzy, Wassif, and Moussa suggest carrying out anomaly detection throughout the various phases of DevOps by linking evidence from the pre-production pipeline with operational monitoring instead of limiting intelligence to alerts that occur after deployment [8]. This kind of architecture is significant since it views reliability as a continuous process; a failed build, odd test behaviour, a risky configuration, and a runtime anomaly are different ways of observing the same delivery system.

Research into incident diagnosis supports this idea. By using the relationships between service calls and abnormal metrics, MicroHECL is able to identify the root causes in large microservice systems, showing how operational diagnosis can be speeded up when the system's topology and telemetry data are analysed together [9]. PatternMatcher also detects root-cause metrics by means of anomaly filtering, pattern classification, and ranking, and achieves high accuracy in the top three predictions on real incident datasets [10]. Although these studies are not CI/CD engines, they do provide the inference layer needed for reliable automated delivery: the pipeline requires machine-readable evidence to establish whether the target

environment is stable enough to accept a change and, following a release, whether any new symptoms are causally linked to the deployment.

The conclusion is that AI should be used alongside rather than taking the place of deterministic pipeline controls. Reproducibility, artifact immutability, branch protection, infrastructure-as-code, and policy checks should continue to be treated as clear rules. AI can be of assistance in cases where uncertainty has real consequences, for example by estimating the probability of failure, detecting abnormal behaviour, prioritising evidence, or choosing a cautious approach to rolling out changes. This kind of separation lowers the risk of an opaque model altering the fundamental safety features of the delivery system and also results in a more clear audit trail since the teams can tell the difference between fixed controls and probabilistic recommendations.

4.2 Observability as the Data Foundation for AI-Enabled Reliability

There is no delivery system powered by AI that can be more reliable than the telemetry it generates. It has been shown by application monitoring tools that dynamic analysis and performance observation are essential for understanding how software behaves [11]. Research into cloud monitoring also indicates that monitoring involves operational costs and therefore must be carefully designed since the act of observing can itself have an impact on response times and on the way resources behave [12]. In the case of CI/CD, this means that the design of the telemetry should start before an AI model is chosen. The pipeline has to correlate together the commits, the build identifiers, the test runs, the artifacts, the deployment versions,

the changes to the infrastructure, the logs, the metrics, the traces, and the outcomes of any incidents using a set of consistent identifiers.

A reliable dataset that is useful must include both signals obtained before and after a deployment. The pre-deployment signals are things such as the size of the change, the services that were modified, updates to dependencies, the results from static analysis, vulnerability findings, test coverage, the history of flaky tests, previous rollback patterns, and environment drift. The post-deployment signals consist of latency percentiles, error rates, saturation, queue depth, the frequency of restarts, trace anomalies, the amount of resource used towards service-level objectives, and customer-impact indicators. The connection between these two sets of signals enables models to learn which features of a change are linked to operational degradation. If that link is missing, then organizations might develop anomaly detectors that are technically accurate but

unable to respond to the operational question of whether this release caused the problem.

Recent research into observability in the context of microservices highlights the difficulty of combining metrics, logs, and traces in distributed deployment environments [29]. This underlines the importance of having an enterprise telemetry contract: each production service should release a minimum number of reliability signals, each deployment should be given tags, and every automated action should be recorded. This standardisation also has a further advantage when it comes to AI: it decreases inconsistency of features among teams, which in turn makes it easier to compare and govern model behaviour. Saudi companies that are adopting AI-enabled DevOps should therefore consider observability not as a dashboard initiative but as a shared data infrastructure for operational assurance.




Figure 1: AI-enabled reliability architecture for automated cloud delivery

4.3 AI-Assisted Testing and Continuous Verification

Testing is the most advanced stage at which to introduce AI into CI/CD since the effects of a wrong recommendation can still be avoided before production takes place. Research involving language models has looked at the problem of flaky tests, the generation of unit tests, test effectiveness, and a number of other software testing tasks. Flakify shows that language representations can be used to predict flaky tests, which are a long-standing cause of false

pipeline failures and the wasting of engineering effort [13]. A large-scale evaluation of language models for the automatic generation of unit tests indicates that the generated tests can achieve substantial coverage, although it also highlights the need to verify compilation, execution, assertions, and fault-detection quality rather than assuming that the generated code is correct by default [14].

A thorough review of large language models in the area of software testing shows that there is

considerable work being done on test preparation, program repair, and other related testing tasks, but it also points out still unresolved issues concerning evaluation, datasets, and generalization [15]. As a possible solution, mutation-informed test generation operates by checking whether the tests actually pick up the faults that have been injected, rather than depending solely on coverage [16]. Recent evaluation studies have made a similar argument, advocating for the use of realistic software and reproducible bug scenarios when assessing language models for testing purposes [17]. Taken together, these findings indicate that AI can play a limited role in continuous verification: it can be used to generate candidate tests, prioritise regression suites, identify unstable tests, and suggest further verification, but it must have objective evidence from actual execution before the release gate can be opened.

The difference is important when it comes to reliable cloud delivery; a model might generate tests that are syntactically valid and either reproduce current behaviour or contain an incorrect assumption. Pipeline design must therefore decouple the generation of tests from authoritative verification. The artifacts that are generated should be run in separate environments, checked to ensure they are deterministic and include meaningful assertions, and high-risk components in particular should be subject to policy. AI can make use of code-change impact analysis in order to select the tests, but any tests that are skipped must be justified by clear evidence relating to dependencies and should periodically be rechecked by running the full suite of tests. Reliability is enhanced when intelligence helps to reduce the cost of unnecessary verification without lowering the evidentiary standard required for deployment.

4.4 Predictive Deployment, Progressive Delivery and Automated Recovery

Generative AI and predictive models go beyond testing by analysing configuration information, change history, and operational data. Research aimed at practical application regards generative AI as a productivity and decision-support tool in the field of software engineering and at the same time stresses the importance of verification and its responsible use [18]. More extensive studies on the application of large language models to software engineering indicate that their value can be found in a number of activities throughout the lifecycle although the results are still dependent on prompt design, the data used, the context and the method of evaluation [19, 20]. A reliable delivery architecture must therefore treat model outputs as evidence when preparing release plans rather than permit free-form generation to directly control production.

Deployment-risk scoring is a particularly reasonable application. A model is able to combine information on the complexity of a change, the importance of the component, test results, the health of the dependencies, associations with past incidents, and the current runtime signals in order to produce a calibrated estimate of the risk. Changes that are low risk can then move automatically through a canary deployment or a small portion of traffic, whereas high-risk changes need extra tests, a broader review, or to be deployed during a safer time window. What is essential here is not the score itself but the action boundary linked to it. In reliability engineering it is necessary to have clear thresholds, override rules, and fallback arrangements when the model is not available or when it is uncertain.

Progressive delivery transforms the boundaries into operational containment. By using canary releases, blue-green deployments, feature flags, and staged traffic routing, the extent of any impact is reduced since only a small group of users is exposed until data has been collected. The pipeline must compare the service indicators specific to each release against a baseline and halt the promotion if errors, latency, or saturation go beyond the agreed-upon limits. An automated rollback is suitable when the failure signal is clear and the way back is known; in the case of unclear degradation, promotion should be paused and passed on to the operators. Thus, a hierarchy of automation is established rather than having to make a simple choice between manual and autonomous deployments.

The research into trustworthy AI for software engineering emphasizes that human and machine capabilities should be designed to supplement each other [21]. This principle applies directly to the way systems are deployed. Humans are good at assessing new situations, understanding business implications, and taking responsibility; machines, on the other hand, are good at reviewing large amounts of historical data and identifying weak patterns. A well-governed CI/CD system makes use of AI to condense the evidence and suggest actions while still keeping human control over exceptional, irreversible, regulated, or high-impact changes. The outcome is that routine deliveries can be made more quickly without normalising the risk of unreviewed changes.

4.5 Continuous Learning, MLOps and Pipeline Governance

By incorporating AI into DevOps, a second lifecycle is introduced within the overall delivery lifecycle: the models used to make pipeline decisions must themselves be developed, validated, deployed, monitored, and retrained. Research into the continuous development of AI models indicates that

data handling, model learning, software development, and system operations constitute an interdependent pipeline, each having its own triggers and quality issues [22]. An enterprise which automates its release decisions without governing the release-risk model ends up with hidden operational debt. Model versions, training data, features, thresholds, evaluation results, and approval history should thus be managed with the same care as application artifacts.

Provenance is the foundation of governance. All AI recommendations must be reproducible using a specific model version and a clearly defined evidence snapshot; this is crucial when teams want to find out why a deployment was approved or rejected. Just as important is drift monitoring since the repository structure, the cloud architecture, the ways in which teams work, and the patterns of workloads all change over time. A risk model developed on the basis of incidents from last year could gradually become less discriminating after a migration or an architectural redesign. Shadow evaluation, regular recalibration, and comparison with simple baselines can show whether a sophisticated model still has a valid reason for carrying out its operational function.

Research into requirements in the field of DevOps also indicates that continuous delivery relies on traceability, communication, non-functional requirements, and integration within the toolchain [23]. Reliability requirements should therefore, where possible, be expressed as machine-readable service objectives and deployment policies. Rather than making a general claim that a service must be highly available, the pipeline can enforce minimum redundancy, test the critical dependencies, check the configuration, and assess the post-deployment indicators. Declarative cloud automation, including Kubernetes-based zero-touch methods, offers a reproducible foundation for this control model [24]. AI becomes safer when it makes recommendations within an environment in which both the desired state and the constraints are already clearly defined.

4.6 Saudi Enterprise Context: Trust, Security and Organizational Readiness

There is no way of transferring international technical evidence to Saudi enterprises without taking the adoption conditions into account. Empirical research carried out on cloud computing in Saudi Arabia shows that trust in cloud service providers is a major factor affecting adoption attitudes, the attitudes being influenced by perceptions concerning the provider's reputation and service-level verification [25]. In the case of AI-enabled CI/CD, this means that the reliability architecture must provide verifiable evidence to both

internal teams and to the relationships with cloud service providers. Deployment histories, service objectives, incident records, access controls, and recovery performance can transform trust from an abstract expectation into an operationally testable feature.

Security and reliability are also closely linked. Research carried out by Saudi scientists into secure, data-intensive edge-cloud workflows indicates that it is necessary to balance quality-of-service goals with threat detection, fault tolerance, and resource efficiency [26]. A CI/CD pipeline which speeds up the deployment of insecure or misconfigured software is not reliable, even if the availability figures appear to stay normal. Security testing, verification of secrets, checks on images and dependencies, control of identities, and the enforcement of policy should therefore be treated as top-tier release gates. Although AI might be used to prioritise alerts or identify unusual patterns, explicit security policy must still be deterministic when it comes to compliance or access boundaries.

Organizational readiness is probably the main obstacle to the introduction of advanced automation. Evidence from Saudi Arabia on the adoption of AI and analytics shows that technology readiness, leadership support, organizational culture, data management, cross-functional integration, and compliance [27], are all relevant to the degree of DevOps maturity. An organization which has fragmented telemetry, lacks definite service ownership, has weak incident learning, and carries out inconsistent deployments cannot safely proceed to autonomous decisions regarding releases. The right approach should be to first standardize delivery and observability, establish service ownership and reliability objectives, automate deterministic controls, introduce AI in an advisory role, and only then grant automated authority after it has been demonstrated that performance is stable.

The situation in Saudi Arabia also supports the idea of human accountability. Since enterprise services can have financial, operational, or public-facing effects, organizational governance may demand that there be clear responsibility for changes involving production. Human-in-the-loop design should not be seen as a temporary flaw on the way to complete automation; for systems with a major impact it can in fact be a permanent control in which the AI carries out triage, prediction, and evidence synthesis while approved positions still have the final say on exceptional releases. This approach combines reliability, governance, and workforce development rather than requiring organizations to decide between automation and accountability.

Table 2: Implementation priorities and measurable controls for Saudi enterprises

Domain	Recommended control	Reliability metric	Saudi-context rationale
Telemetry	Correlate commit, artifact, deployment, metrics, logs and traces	Trace completeness; signal coverage	Creates auditable service evidence and supports trusted operations
Verification	Risk-based testing with objective execution gates	Escaped defects; flaky-test rate; test recall	Preserves quality while reducing unnecessary pipeline delay
Deployment	Progressive release with automated stop and rollback	Change failure rate; rollback correctness	Limits customer and operational exposure during frequent change
AIOps	Anomaly detection and root-cause support	Detection time; diagnosis time; alert precision	Supports faster recovery in complex cloud and microservice estates
Governance	Policy-as-code, model provenance, approval tiers	Override rate; audit completeness	Aligns automation with service-level, security and accountability needs
Readiness	Phased authority based on data and team maturity	Automation success; recovery time; training coverage	Allows capability to grow with organizational trust and skills

5. Review-Derived Reference Model

The synthesis presents a six-stage reference model for the reliable cloud delivery enabled by AI. In the first stage, change capture involves versioning commits, configuration modifications, dependency updates and infrastructure changes and linking them to service ownership. In the second stage, automated verification generates baseline evidence through builds, static analysis, security checks and deterministic tests. In the third stage, intelligent assessment carries out the prioritisation of tests, detects any unusual behaviour in the pipeline, estimates the risk of a release and identifies any gaps in the evidence. In the fourth stage, governed release converts risk into an approved approach such as direct deployment, canary, blue-green, additional review or rejection. In the fifth stage, runtime verification determines whether the release remains within accepted limits by means of correlated metrics, logs, traces and service objectives. In the sixth stage, learning takes place when successful deployments, incidents, rollbacks, overrides and the findings from postmortems are used to update the future rules and the model datasets.

The model intentionally keeps prediction and enforcement separate. Although it can calculate a

risk score, it is policy that determines the score at which action is taken. Reliability models should likewise be evaluated quantitatively and not simply trusted since they involve the use of sophisticated algorithms. Research into the reliability and resilience of machine learning shows the benefits of monitoring degradation and recovery over time [28]. When this approach is applied to CI/CD, it means that, together with the usual DevOps metrics, model precision, false-safe decisions, false blocks, calibration, drift, and operational consequences should all be monitored.

Figure 1 shows the architecture and Figure 2 illustrates the closed feedback loop. Both of these diagrams highlight the same idea: automation can only be reliable if information is sent from production back to the delivery system. A pipeline that terminates at deployment has no way of finding out whether its tests, risk assessments, or policies were effective. On the other hand, an operations platform which is unable to associate incidents with the specific change and release decision cannot improve its ability to prevent such incidents. It is therefore necessary structurally to have end-to-end traceability if AI assistance is to be linked with reliable service delivery.




Figure 2: Closed-loop reliability control model for Saudi enterprise cloud services

6. DISCUSSION

The literature does support AI-enabled DevOps, but it does not support indiscriminate autonomy. The most compelling evidence relates to bounded tasks that have observable outcomes, such as identifying anomalies, ranking the most probable root causes, predicting flaky tests, generating candidate tests, and assessing software artifacts. Research into transfer learning also indicates that learned software representations can be reused for different code-related tasks, although the degree of effectiveness varies depending on the specific task and data characteristics [30]. These capabilities can increase both the speed and the scope of evidence processing, but they only become enterprise reliability mechanisms when they are linked to explicit release policies and checked against actual production results.

This results in a practical difference between automation maturity and autonomy maturity. Automation maturity refers to things such as reproducible builds, the execution of tests, deployment, the provisioning of infrastructure, and monitoring. Autonomy maturity relates to whether an intelligent component is able to select, modify, or halt those automated actions. Organizations must not move on to the second aspect faster than the first. A team which is not capable of reliably reproducing a deployment should not permit a model to optimize its deployment path. In a similar way, a team that does not have clearly defined service objectives should not

be able to judge whether an AI-generated rollback decision has improved reliability.

The review also states that general DevOps metrics should be supplemented by those specific to AI. While figures such as change failure rate, deployment frequency, lead time, recovery time, availability, latency, and service-level compliance are still important, they should be combined with decision metrics including risk-score calibration, alert precision, test-selection recall, automated rollback correctness, the percentage of model recommendations that are overridden, and the amount of time spent waiting for human approval. It is these measures that show whether AI eliminates useful friction or simply shifts the effort into the area of model supervision and incident investigation.

Saudi companies would do better with a step-by-step approach to operation. In the first stage the pilots could concentrate on low-risk recommendations, for example those relating to predicting flaky tests, clustering logs, or prioritising tests. In the following stage AI could be used to suggest deployment strategies while human beings approved the promotion to production. In more developed environments the system could automate low-risk releases using a gradual exposure and objective rollback mechanisms. High-impact services should continue to require mandatory approval and maintain strict evidence requirements forever. This kind of progression according to risk enables organisations to build up their capabilities and

confidence while at the same time producing tangible evidence of reliability.

Reviews have certain limitations since the thirty studies involve a wide variety of systems, datasets, and outcomes, and a number of AI techniques are assessed in controlled or task-specific environments rather than throughout the full enterprise delivery chain. Since generative models develop more quickly than the usual publication cycles, the evidence regarding particular model families can become outdated rapidly. Research specifically concerned with AI-controlled CI/CD in Saudi Arabia is still limited; for this reason the national interpretation draws on both local evidence of adoption and on international technical studies. The reference model that is proposed therefore ought to be seen as a synthesis to aid implementation and for empirical testing, not as proof that a particular technology stack will lead to a fixed reliability gain.

7. Research Gaps and Future Agenda

There are five areas that should be given priority. First, end-to-end studies ought to assess AI-assisted delivery from the time of commit up to production recovery rather than simply provide accounts of individual model accuracies. Second, research should make use of shared datasets that connect code changes, pipeline events, deployment strategies, telemetry, and the results of incidents all without revealing sensitive enterprise information. Third, deployment-risk models need calibration research in order that the probabilities they give match the actual failure rates experienced in operation across different service classes. Fourth, organisations should have methods for measuring the contribution to reliability of human-AI decision patterns, such as in cases where overrides avoid failures and when excessive approval undermines the benefits of automation. Fifth, the studies carried out in Saudi enterprises should examine these mechanisms within the context of local cloud architectures, governance practices, data-location requirements, workforce structures, and service expectations.

Future studies should involve comparing simple, interpretable models with more complex language or deep-learning methods when using the same operational data. The research should also look at the various ways in which the AI control plane itself might fail, for example by having unavailable inference services, suffering from corrupted telemetry, being given adversarial inputs, receiving stale features, or being presented with misleading model confidence. A reliable enterprise platform has to be able to degrade safely when intelligence is not available. Moreover, the field needs more rigorous economic analysis that links AI-enabled CI/CD to the

costs avoided from outages, the engineering effort involved, cloud expenditure, and the opportunities for releases. Only in this way will organisations be able to tell the difference between automation that is technically impressive and changes that actually lead to an improvement in service reliability.

8. CONCLUSION

AI can be used to enhance DevOps and automated CI/CD by improving the quality, speed, and continuity of reliability evidence. The most defensible uses of AI are predictive failure assessment, intelligent verification, anomaly detection, root-cause support, progressive-release decision support, and feedback-driven recovery. The value of these applications relies on observability, traceability, deterministic policy, reversible deployment, model governance, and human accountability. For Saudi enterprises, aspects such as cloud trust, secure service operation, organizational readiness, and measurable service-level assurance are just as important as model capability. The architecture derived from the review therefore positions AI within a bounded closed loop rather than placing it above the delivery system. Organizations which first standardise their telemetry and delivery processes and then gradually introduce intelligence can achieve faster release cycles without viewing automation as a replacement for engineering control. The following step in research should be production-scale validation that links model decisions directly to service outcomes, incident reduction, recovery performance, and enterprise value.

REFERENCES

1. Barenkamp M, Rebstadt J, Thomas O. Applications of AI in classical software engineering. *AI Perspectives*. 2020;2:1. doi:10.1186/s42467-020-00005-4.
2. Martinez-Fernandez S, Bogner J, Franch X, Oriol M, Siebert J, Trendowicz A, Vollmer AM, Wagner S. Software Engineering for AI-Based Systems: A Survey. *ACM Transactions on Software Engineering and Methodology*. 2022;31(2):37. doi:10.1145/3487043.
3. Watson C, Cooper N, Nader Palacio D, Moran K, Poshyvanyk D. A Systematic Literature Review on the Use of Deep Learning in Software Engineering Research. *ACM Transactions on Software Engineering and Methodology*. 2022;31(2):32. doi:10.1145/3485275.
4. Kotti Z, Galanopoulou R, Spinellis D. Machine Learning for Software Engineering: A Tertiary Study. *ACM Computing Surveys*. 2023;55(12):256. doi:10.1145/3572905.
5. Wen L, Qian H, Liu W. Research on Intelligent Cloud Native Architecture and Key Technologies Based on DevOps Concept. *Procedia Computer*

- Science. 2022;208:590-597. doi:10.1016/j.procs.2022.10.082.
6. Mowad AM, Fawareh H, Hassan MA. Effect of Using Continuous Integration (CI) and Continuous Delivery (CD) Deployment in DevOps to reduce the Gap between Developer and Operation. 2022 International Arab Conference on Information Technology. 2022. doi:10.1109/ACIT57182.2022.9994139.
7. Beetz F, Harrer S. GitOps: The Evolution of DevOps? IEEE Software. 2022;39(4):70-75. doi:10.1109/MS.2021.3119106.
8. Fawzy AH, Wassif K, Moussa H. Framework for automatic detection of anomalies in DevOps. Journal of King Saud University - Computer and Information Sciences. 2023;35:8-19. doi:10.1016/j.jksuci.2023.02.010.
9. Liu D, He C, Peng X, Lin F, Zhang C, Gong S, Li Z, Ou J, Wu Z. MicroHECL: High-Efficient Root Cause Localization in Large-Scale Microservice Systems. IEEE/ACM ICSE-SEIP. 2021:338-347. doi:10.1109/ICSE-SEIP52600.2021.00043.
10. Wu C, Zhao N, Wang L, Yang X, et al. Identifying Root-Cause Metrics for Incident Diagnosis in Online Service Systems. IEEE ISSRE. 2021. doi:10.1109/ISSRE52982.2021.00022.
11. Hasselbring W, van Hoorn A. Kieker: A monitoring framework for software engineering research. Software Impacts. 2020;5:100019. doi:10.1016/j.simpa.2020.100019.
12. Khalil MM, Khomonenko AD, Matushko MD. Measuring the effect of monitoring on a cloud computing system by estimating the delay time of requests. Journal of King Saud University - Computer and Information Sciences. 2022;34:3968-3972. doi:10.1016/j.jksuci.2021.02.001.
13. Fatima S, Ghaleb TA, Briand L. Flakify: A Black-Box, Language Model-Based Predictor for Flaky Tests. IEEE Transactions on Software Engineering. 2023;49(4):1912-1927. doi:10.1109/TSE.2022.3201209.
14. Schafer M, Nadi S, Eghbali A, Tip F. An Empirical Evaluation of Using Large Language Models for Automated Unit Test Generation. IEEE Transactions on Software Engineering. 2024;50(1):85-105. doi:10.1109/TSE.2023.3334955.
15. Wang J, Huang Y, Chen C, Liu Z, Wang S, Wang Q. Software Testing With Large Language Models: Survey, Landscape, and Vision. IEEE Transactions on Software Engineering. 2024;50(4):911-936. doi:10.1109/TSE.2024.3368208.
16. Moradi Dakhel A, Nikanjam A, Majdinasab V, Khomh F, Desmarais MC. Effective test generation using pre-trained Large Language Models and mutation testing. Information and Software Technology. 2024;171:107468. doi:10.1016/j.infsof.2024.107468.
17. Li Y, Liu P, Wang H, Chu J, Wong WE. Evaluating large language models for software testing. Computer Standards & Interfaces. 2025;93:103942. doi:10.1016/j.csi.2024.103942.
18. Ebert C, Louridas P. Generative AI for Software Practitioners. IEEE Software. 2023;40(4):30-38. doi:10.1109/MS.2023.3265877.
19. Fan A, Gokkaya B, Harman M, Lyubarskiy M, Sengupta S, Yoo S, Zhang JM. Large Language Models for Software Engineering: Survey and Open Problems. IEEE/ACM ICSE-FoSE. 2023:31-53. doi:10.1109/ICSE-FoSE59343.2023.00008.
20. Hou X, Zhao Y, Liu Y, Yang Z, Wang K, Li L, Luo X, Lo D, Grundy J, Wang H. Large Language Models for Software Engineering: A Systematic Literature Review. ACM Transactions on Software Engineering and Methodology. 2024;33(8):220. doi:10.1145/3695988.
21. Lo D. Trustworthy and Synergistic Artificial Intelligence for Software Engineering: Vision and Roadmaps. IEEE/ACM ICSE-FoSE. 2023. doi:10.1109/ICSE-FoSE59343.2023.00010.
22. Steidl M, Felderer M, Ramler R. The pipeline for the continuous development of artificial intelligence models: Current state of research and practice. Journal of Systems and Software. 2023;199:111615. doi:10.1016/j.jss.2023.111615.
23. Hernandez R, Moros B, Nicolas J. Requirements management in DevOps environments: a multivocal mapping study. Requirements Engineering. 2023;28:317-346. doi:10.1007/s00766-023-00396-w.
24. Mamatas L, Skaperas S, Sakellariou I. ClusterSlice: Slicing resources for zero-touch Kubernetes-based experimentation. Future Generation Computer Systems. 2024;161:1-10. doi:10.1016/j.future.2024.06.038.
25. Alghaith W. The Role of Trust in Cloud Service Providers on the Adoption of Cloud Computing in Saudi Arabia: An Empirical Investigation. Journal of Computer Science. 2023;19(9):1107-1123. doi:10.3844/jcssp.2023.1107.1123.
26. Aljohani M. Transforming data-intensive workflows: A cutting-edge multi-layer security and quality aware security framework. Concurrency and Computation: Practice and Experience. 2024;36(13):e8068. doi:10.1002/cpe.8068.
27. Alzahrani S. Assessing artificial intelligence and advanced analytics adoption: A technological, organizational, and environmental perspective. Edelweiss Applied Science and Technology. 2025;9(7):312-325. doi:10.55214/25768484.v9i7.8567.

28. Faddi Z, da Mata K, Silva P, Nagaraju V, Ghosh S, Kul G, Fiondella L. Quantitative assessment of machine learning reliability and resilience. Risk Analysis. 2025;45(4):790-807. doi:10.1111/risa.14666.
29. Faseeha U, Syed HJ, Samad F, Zehra S, Ahmed H. Observability in Microservices: An In-Depth Exploration of Frameworks, Challenges, and Deployment Paradigms. IEEE Access. 2025;13:72011-72039. doi:10.1109/ACCESS.2025.3562125.
30. Mastropaolo A, Cooper N, Nader Palacio D, Scalabrino S, Poshyvanyk D, Oliveto R, Bavota G. Using Transfer Learning for Code-Related Tasks. IEEE Transactions on Software Engineering. 2022. doi:10.1109/TSE.2022.3183297.