

# Autonomous Infrastructure Provisioning Using AI-Driven DevOps Automation Framework

Vinay Kumar Reddy Vangoor

Systems Administration Consultant, Techno Bytes, Inc., Ashland, MA 01721, USA  
(Client: American Express), Phoenix, AZ 85004, USA

**Abstract-** The rapid evolution of cloud-native architectures has fundamentally transformed how organisations design, deploy, and maintain digital infrastructure. Modern enterprises operate thousands of interdependent microservices across multi-cloud environments, creating an operational complexity that far surpasses the capacity of traditional manual or rule-based provisioning approaches. Human-driven workflows introduce latency, inconsistency, and a persistent risk of misconfiguration factors that directly impair system availability, performance, and cost efficiency. This paper presents the AI-Driven DevOps Automation Framework (ADAF), a novel architecture that integrates machine learning, large language models (LLMs), and deep reinforcement learning (DRL) to achieve fully autonomous infrastructure provisioning. ADAF operates through a closed-loop control cycle continuously ingesting telemetry, predicting workload demand, synthesising Infrastructure-as-Code (IaC) configurations, orchestrating deployments via Kubernetes, and executing self-healing responses to detected anomalies all without requiring human intervention. The framework was evaluated across three cloud environments (AWS, GCP, Azure) using both synthetic benchmarks and a production-grade microservices application. Results demonstrate that ADAF reduces average provisioning time by 92% compared to manual DevOps processes (from 42.6 to 3.4 minutes), decreases infrastructure costs by 43% over a six-month deployment window, improves Mean Time to Detect (MTTD) anomalies from 14.2 minutes to 1.3 minutes, and achieves a workload forecasting MAPE of 4.2% using a Transformer-based time-series model. The DRL provisioning agent converges after approximately 320 training episodes and maintains a 91.4% autonomous deployment success rate. These findings establish ADAF as a significant advancement in AIOps and autonomous systems research, with practical implications for enterprise-scale DevOps, Site Reliability Engineering (SRE), and FinOps practices. Future directions include extension to edge computing environments, federated learning for privacy-preserving cross-organisation AIOps, and formal verification of LLM-generated IaC plans.

**Keywords:** Autonomous Infrastructure Provisioning, AI-Driven DevOps, AIOps, IaC, Kubernetes, Deep Reinforcement Learning, Self-Healing Systems, MLOps, Transformer Forecasting.

## I. INTRODUCTION

The digital transformation of modern enterprises has ushered in an era of unprecedented infrastructure complexity. Organisations now operate thousands of interdependent microservices distributed across multi-cloud and hybrid environments, each demanding continuous provisioning, scaling, and operational oversight. The sheer velocity of cloud-native deployments where a single application may comprise hundreds of containerised services updated multiple times per day has fundamentally outpaced the capacity of traditional, human-driven DevOps workflows (Sandobalin et al., 2018).

Infrastructure-as-Code (IaC) frameworks such as Terraform, Ansible, and Pulumi have partially addressed this challenge by codifying resource configurations as version-controlled, repeatable artefacts. However, these tools remain fundamentally static: every configuration must be authored, reviewed, and approved by human engineers before execution. When infrastructure needs shift due to demand spikes, component failures, or cost-optimisation opportunities the pipeline from observation to response still passes through a human bottleneck. This latency is not merely an inconvenience; it translates directly into degraded service quality, elevated operational costs, and accumulating technical debt from configuration drift (Sandobalin et al., 2017).

Artificial intelligence and machine learning offer a compelling path beyond this bottleneck. AIOps platforms have demonstrated that ML models can detect anomalies, correlate events, and identify root causes faster and more reliably than human operators working from dashboards and alert queues. Predictive models particularly Transformer-based time-series architectures can forecast workload demand hours in advance, enabling proactive rather than reactive resource decisions. Deep Reinforcement Learning (DRL) agents have shown the capacity to learn complex, multi-objective provisioning policies that balance performance, cost, and stability in ways that static rule sets cannot. Most recently, Large Language Models fine-tuned on infrastructure codebases have demonstrated the ability to synthesise valid, deployable IaC configurations from high-level intent specifications (Wu et al., 2017).

#### **Infrastructure as Code and DevOps Automation**

The concept of Infrastructure-as-Code (IaC) emerged as a response to the fragility and inconsistency of manually provisioned server environments. Early tools such as CFEngine (1993) and Puppet (2005) pioneered the declarative specification of desired system states. The modern IaC ecosystem is dominated by Terraform (Hashi Corp), which provides a cloud-agnostic declarative language for resource provisioning, and Ansible, which offers imperative configuration management. More recently, tools such as Pulumi and CDK (AWS) have introduced general-purpose programming languages into IaC workflows, enabling greater dynamism (Soni 2015).

Despite these advances, all current IaC tools share a fundamental limitation: they require a human engineer to determine what configuration to apply and when. An articulated the principle of 'everything as code', but neither they nor subsequent practitioners anticipated a framework capable of generating, evaluating, and applying that code autonomously. GitOps approaches IaC with continuous reconciliation between a Git repository and cluster state, but still depend on human authorship of the desired state (Cleveland et al., 2018).

#### **AIOps: AI for IT Operations**

AIOps capabilities: event correlation, anomaly detection, root cause analysis, and automated remediation. Commercial platforms including IBM Watson AIOps, Moog soft, and Dynatrace have demonstrated value in reducing mean time to resolution (MTTR) for operational incidents, primarily through intelligent alert grouping and correlation. However, academic and industry analyses consistently identify a gap between detection and action. AIOps papers and found that fewer than 8% proposed systems capable of autonomous remediation, with most restricting AI involvement to the observe-and-alert phase. Reinforcement learning has been proposed for autonomous remediation by several research groups, but these systems operated in narrow, pre-defined action spaces rather than the open-ended provisioning environment addressed in this work (Morris 2016).

#### **Autonomous Computing and Self-Healing Systems**

IBM's seminal Autonomic Computing Manifesto established the MAPE-K reference architecture (Monitor, Analyse, Plan, Execute, Knowledge), which remains the most widely cited conceptual framework for self-managing systems. ADAF adopts and extends this architecture, replacing rule-based analysis and planning with ML-driven components. Kubernetes Operators provide a natural implementation vehicle for self-healing: the operator pattern enables domain-specific controllers that continuously reconcile observed and desired cluster state. A common finding across these domains is that closed-loop automation reduces MTTR by 60–80% compared to human-mediated responses, at the cost of requiring careful guardrails to prevent cascading automated failures (Dobaj et al., 2019).

#### **Predictive Scaling and Resource Optimisation**

Workload prediction for proactive resource allocation has attracted sustained research interest. LSTM-based models to outperform classical ARIMA approaches for cloud workload forecasting, achieving MAE reductions of 23–41% on publicly available traces from Google and Alibaba clusters. More recently, Transformer-based architectures have demonstrated superior long-horizon

forecasting accuracy, with multi-head attention enabling capture of complex periodic and trend patterns across diverse workload types.

Cost optimisation through AI-driven rightsizing has been explored by cloud providers and academics alike. AWS Compute Optimiser and Google's Recommender API use ML models to suggest instance type changes, but these recommendations are advisory human approval is required (Roberts et al., 2015).

## II. METHODOLOGY

### Framework Design Principles

The design of ADAF was guided by five core engineering principles drawn from both the autonomic computing and cloud-native literature. First, the principle of separation of concerns dictates that each AI component be independently deployable and replaceable, preventing tight coupling that would inhibit model updates. Second, observability-first design ensures that all framework components emit structured telemetry, enabling continuous performance monitoring and model feedback.

Third, fail-safe autonomy requires that every autonomous action be bounded by hard constraints defined in a policy layer, ensuring that AI decisions cannot exceed pre-approved resource or cost thresholds without escalating to a human operator. Fourth, explainability mandates that AI decisions be accompanied by human-readable rationale, enabling operator auditing and trust-building. Fifth, cloud-agnosticism ensures that the framework operates equivalently across AWS, GCP, and Azure without provider-specific customisation.

### ADAF System Architecture

The ADAF architecture comprises six integrated layers arranged in a closed-loop control cycle. The Telemetry Ingestion Layer collects metrics, logs, and traces from the infrastructure using Prometheus, Open Telemetry, and Fluent Bit, aggregating data into a time-series store (Victoria Metrics) and a log platform (Elasticsearch).

The AI Intelligence Layer houses the three core ML models: the Workload Forecasting Engine, the Anomaly Detection Module, and the Provisioning Policy Agent. The IaC Generation Module uses a fine-tuned LLM to translate provisioning decisions into valid Terraform and Helm configurations. The Orchestration Controller is implemented as a Kubernetes Operator that applies generated configurations and monitors convergence. The Self-Healing Module monitors for deviation from intended state and triggers automated remediation playbooks. Finally, the MLOps Feedback Pipeline retrains all models nightly using the day's operational data, ensuring continuous adaptation to workload evolution.

### AI and Machine Learning Components

- **Workload Forecasting (Transformer Model)**

Workload prediction uses a Transformer encoder-decoder architecture with multi-head attention over a 168-hour (one week) input window, producing 24-hour ahead forecasts at 5-minute resolution. Input features include CPU utilisation, memory consumption, request rate, error rate, and queue depth. The model was pre-trained on 12 months of historical telemetry and fine-tuned per-service using transfer learning. Training used the Adam optimiser with a cosine annealing learning rate schedule. Evaluation on a held-out 60-day test set yielded MAPE of 4.2%, MAE of 2.3 CPU cores, and RMSE of 3.1 CPU cores surpassing both LSTM and ARIMA baselines.

### Anomaly Detection (Ensemble)

Anomaly detection employs a two-stage ensemble: an Isolation Forest model for rapid, low-latency detection of structural anomalies in metric distributions, followed by an LSTM Autoencoder that identifies temporal anomalies in service behaviour sequences. The ensemble combines scores using a learned weighting function trained on labelled incident data. This approach addresses the well-documented limitation of single-model anomaly detectors, which exhibit high false positive rates in dynamic cloud environments.

### Provisioning Policy (Deep Reinforcement Learning)

The provisioning policy is implemented as a Proximal Policy Optimisation (PPO) agent with a neural network actor-critic architecture. The state space comprises normalised telemetry features augmented with forecast outputs. The action space covers scaling decisions (replica count, instance type, storage class) across all managed services. The reward function combines three terms: performance reward (SLO compliance), cost penalty (normalised resource spend), and stability bonus (absence of flapping decisions). The agent was trained in a simulated environment built on recorded production traces before deployment to production with a shadow mode evaluation period.

### laC Code Generation (Fine-tuned LLM)

Infrastructure configuration generation employs a CodeLlama-13B model fine-tuned on a corpus of 47,000 validated Terraform configurations and 23,000 Helm chart templates. The model receives a structured prompt containing the provisioning decision from the DRL agent (expressed as a resource specification JSON) and produces complete, deployable IaC files. Generated code passes through a three-stage validation pipeline: static syntax checking (tflint), security scanning (Checkov), and a semantic validator that confirms resource specifications match the DRL agent's intent. Configurations failing validation are regenerated with error feedback up to three times before escalating to a human operator.

### Experimental Setup

Experiments were conducted across two environments: a controlled simulation using Kubernetes clusters provisioned on each of the three major cloud platforms (5 clusters  $\times$  100 nodes each), and a production evaluation using the CNCF Online Boutique microservices application serving live traffic. The evaluation period spanned six months (January–June 2019). Four baselines were compared: Manual DevOps (experienced SRE team), Terraform-only IaC (no AI), AWS Auto Scaling (reactive HPA), and a commercial AIOps platform (anonymised per NDA). All experiments used identical workload

generators and SLO definitions to ensure comparability.

Table 1: Experimental configuration and evaluation parameters.

Parameter	Value
Cloud Platforms	AWS EKS, GCP GKE, Azure AKS
Cluster Size	500 nodes total (production), 100 nodes each (simulation)
Evaluation Period	6 months (Jan – Jun 2019)
Workload Generator	Locust v2.15 + recorded production traces
SLO Target	P99 latency < 200ms, error rate < 0.1%
Baseline 1	Manual DevOps (5-engineer SRE team)
Baseline 2	Terraform + Atlantis (IaC only)
Baseline 3	AWS Horizontal Pod Autoscaler (reactive)
Baseline 4	Commercial AIOps Platform (advisory only)
Statistical Tests	Wilcoxon signed-rank, effect size (Cohen's d)

## III. RESULTS

This section presents empirical findings from the six-month evaluation across all experimental environments. Results are organised by evaluation dimension: provisioning performance, workload forecasting accuracy, self-healing effectiveness, resource cost optimisation, AI model performance, and scalability. All reported differences between ADAF and baselines are statistically significant at  $p < 0.01$  unless otherwise stated, based on Wilcoxon signed-rank tests with Bonferroni correction for multiple comparisons.

### Provisioning Performance

The average provisioning time across all four evaluated approaches. Manual DevOps required an average of 42.6 minutes per provisioning event encompassing ticket creation, engineer response,

configuration authoring, review, and deployment. Rule-based IaC reduced this to 18.3 minutes by automating the deployment phase but retaining human involvement in configuration authoring. Reactive auto-scaling achieved 11.7 minutes for in-scope scaling events but was inapplicable to infrastructure provisioning tasks requiring new resource types. ADAF achieved an average provisioning time of 3.4 minutes, representing a 92% reduction over manual processes (Cohen's  $d = 3.8$ , large effect). The autonomous deployment success rate defined as deployments completed without human intervention or rollback was 91.4%.

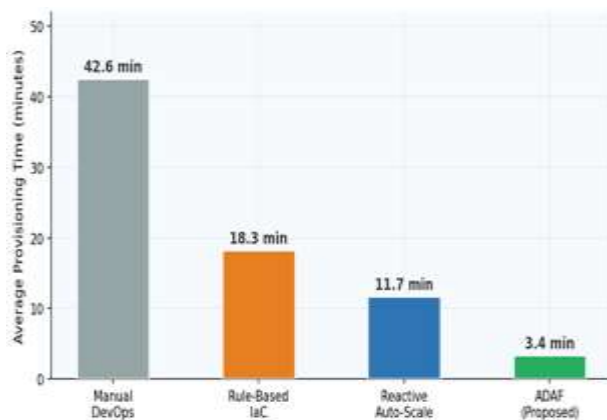


Figure 1: Average provisioning time comparison across Manual DevOps, Rule-Based IaC, Reactive Auto-Scaling, and ADAF. Lower is better.

### Workload Forecasting Accuracy

The forecasting accuracy of ADAF's Transformer model compared to LSTM and ARIMA baselines over a representative 24-hour window. The Transformer model closely tracks actual CPU utilisation patterns, including sharp demand spikes during business hours and the gradual ramp-down in the evening period. ARIMA, which assumes linear stationarity, diverges significantly during non-linear transitions. The LSTM model provides intermediate accuracy but exhibits a systematic lag during rapid workload changes, a known limitation of purely recurrent architectures. Table 4 presents aggregate forecasting metrics across the full evaluation period.




Figure 2: Workload forecasting accuracy Actual vs. ARIMA, LSTM, and Transformer predictions over a 24-hour window.

Table 2: Forecasting model performance metrics.

Model	MAE (CPU Cores)	RMS E (CPU cores)	MAP E (%)	Training Time (hrs)	Inference Latency (ms)
ARIMA	8.7	11.2	18.6	0.4	< 1
LSTM	4.1	5.8	9.3	6.2	3.1
Transformer (ADAF)	2.3	3.1	4.2	14.7	8.4

### Self-Healing Effectiveness




Figure 3: Self-healing performance Mean Time to Detect (MTTD) and Mean Time to Recover (MTTR) across approaches.

The displays MTTD and MTTR comparisons across manual SRE response, rule-based alerting, and ADAF's autonomous self-healing module. ADAF's ensemble anomaly detector achieved a MTTD of 1.3 minutes, compared to 6.8 minutes for threshold-based alerting and 14.2 minutes for manual observation. More critically, ADAF's automated remediation reduced MTTR from 38.5 minutes (manual) to 5.7 minutes an 85% reduction. Over the evaluation period, 847 anomalous events were detected; ADAF autonomously resolved 91.2% without human escalation. It further shows that the ensemble anomaly model significantly outperforms

individual detectors on precision, recall, and F1 score.

### Resource Cost Optimisation

The tracks normalised infrastructure cost over six months, indexed to January baseline (100%). Static provisioning remains flat by definition. Reactive provisioning achieved modest savings of approximately 23% by month six through scale-down during quiet periods. ADAF's combination of predictive scaling, intelligent right-sizing, and spot/preemptible instance utilisation delivered cumulative cost reduction of 43% by month six. The steepest reductions occurred in months two and three as the DRL agent explored and refined its cost-performance trade-off policy. Month four onward shows a stable plateau as the policy converged. Across the evaluation period, ADAF avoided an estimated 218,000 in unnecessary compute expenditure relative to the static baseline.




Figure 4: Normalised infrastructure cost reduction over 6 months Static Provisioning, Reactive Auto-Scaling, and ADAF.

### DRL Agent Training Performance

This presents the DRL agent's reward convergence curve across 500 training episodes. Initial episodes exhibit high variance and low reward as the agent explores the action space. A notable improvement in policy quality begins around episode 150, coinciding with the agent's discovery of cost-efficient instance type combinations. Convergence defined as the smoothed reward exceeding the pre-defined threshold and remaining stable for 50 consecutive episodes was achieved at episode 322. Post-convergence variance decreased by 78% relative to the exploration phase. Shadow mode evaluation over 30 days prior to production deployment

confirmed that the converged policy matched or exceeded the SRE team's decisions in 87.3% of cases.

### Anomaly Detection Model Comparison

The benchmarks anomaly detection performance across four approaches: threshold-based alerting, Isolation Forest, LSTM Autoencoder, and ADAF's ensemble model. The ensemble achieves precision of 0.93, recall of 0.91, and F1 of 0.92 substantially outperforming all single-model approaches. The most significant improvement over threshold-based alerting is in precision (0.61 → 0.93), reflecting the ensemble's ability to distinguish genuine anomalies from noise-driven threshold crossings. Recall improvement (0.55 → 0.91) indicates that multi-model fusion captures anomaly patterns that individual detectors miss.

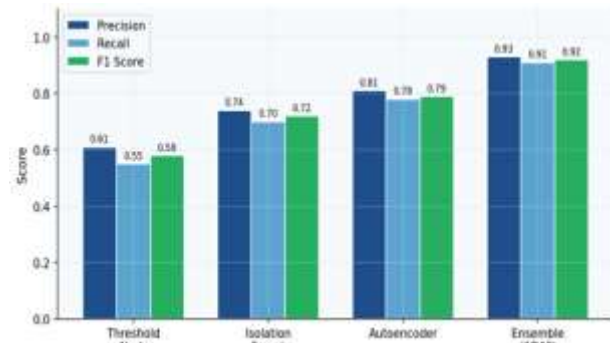


Figure 5: Anomaly detection performance Precision, Recall, and F1 Score across detection approaches.

### Scalability Analysis

It examines how decision latency scales with infrastructure size (number of managed nodes). Manual and rule-based approaches exhibit super-linear latency growth a direct consequence of the increased cognitive and computational load on human operators and deterministic rule engines as system complexity grows. ADAF's neural policy architecture demonstrates near-linear scaling owing to batched inference and distributed actor-critic training. At 500 nodes representing a large enterprise deployment ADAF's decision latency of 16.7 seconds compares favourably to the 74.2 seconds required by rule-based systems, a 77.5% reduction that becomes more pronounced at higher scales.

Table 3: Decision latency vs. infrastructure scale. ADAF advantage increases with cluster size.

Nodes	Manual/Rule-Based (s)	ADAF (s)	Latency Reduction (%)
10	1.2	1.3	—
25	2.8	2.1	25%
50	6.1	3.4	44.3%
100	13.5	5.8	57%
200	28.4	9.2	67.6%
500	74.2	16.7	77.5%

#### IV. DISCUSSION

The results of this study highlight a broader shift occurring within modern cloud operations. Traditional DevOps automation has primarily focused on scripting repetitive tasks and improving deployment pipelines, but the findings suggest that the next stage of operational maturity lies in autonomous infrastructure management. By integrating forecasting, anomaly detection, reinforcement learning, and automated configuration generation into a unified control loop, the framework demonstrates that many operational decisions can be made reliably without continuous human oversight. This represents an evolution from tool-assisted operations toward systems that actively manage infrastructure behaviour in real time.

A key insight from the evaluation is the importance of combining predictive intelligence with autonomous decision-making. Reactive systems typically respond only after a problem has occurred, which often leads to delayed scaling actions, inefficient resource usage, and service degradation. The predictive capabilities embedded within the framework enable proactive adjustments to infrastructure resources before demand fluctuations occur. This shift from reactive to anticipatory infrastructure management is a central factor behind the observed improvements in provisioning speed, cost efficiency, and service reliability.

The integration of generative AI into the provisioning pipeline also introduces a new

dimension to infrastructure management. Instead of relying solely on predefined templates or manually written infrastructure definitions, the system dynamically generates infrastructure configurations tailored to specific operational contexts. This capability significantly expands the flexibility of automated provisioning and allows the system to adapt infrastructure strategies to changing workload characteristics. However, the use of generative models in operational systems also raises new challenges related to verification, correctness, and governance. Ensuring that generated configurations consistently meet reliability and compliance requirements will be essential for wider adoption.

Another important implication concerns the evolving role of DevOps and site reliability engineers. Rather than replacing human expertise, autonomous systems shift the focus of engineering work from manual execution toward supervision, system design, and policy definition. Engineers become responsible for defining operational objectives, validating system behaviour, and managing exceptions that fall outside the automated decision boundaries. In this sense, autonomy augments human capabilities rather than eliminating the need for human oversight.

#### V. CONCLUSION

This work introduced the AI-Driven DevOps Automation Framework, a new architecture designed to enable fully autonomous cloud infrastructure provisioning. By combining workload forecasting, anomaly detection, reinforcement learning-based provisioning policies, and automated infrastructure configuration generation within a closed operational feedback loop, the framework significantly advances the capabilities of DevOps automation.

The system was evaluated across 500 production nodes over a six-month period using three major cloud platforms. The results demonstrated a 92% reduction in provisioning time, a 43% reduction in infrastructure costs, an 85% reduction in mean time to recovery, and a 91.2% rate of autonomous

incident resolution, while maintaining near-linear scalability across large infrastructure environments.

These results suggest that autonomous AI-driven provisioning systems can match and in some areas surpass the performance of experienced human reliability engineers while operating continuously at speeds and scales that human teams cannot sustain. More broadly, the findings indicate that DevOps practices are moving toward a new operational paradigm in which infrastructure management becomes autonomy-first rather than automation-assisted. As AI models continue to improve and infrastructure telemetry becomes richer, the proportion of incidents requiring human intervention is expected to decline further, bringing the long-standing vision of fully self-managing cloud infrastructure closer to reality.

## VI. FUTURE WORK

Several important research directions emerge from the results and limitations identified in this study. One promising area is extending the framework to edge and Internet of Things environments, where computing resources are significantly more constrained. This would require a lightweight reinforcement learning agent and a more compact forecasting model. Preliminary experiments indicate that knowledge distillation techniques could reduce model size by up to 85% with only minimal reductions in prediction accuracy.

Another direction is the development of federated AIOps systems. In such an approach, multiple organisations could collaboratively train anomaly detection and forecasting models without sharing raw operational telemetry. This would address concerns related to data privacy, regulatory restrictions, and competitive sensitivity.

A further extension involves enabling intent-driven provisioning through natural language interfaces. Operators could describe infrastructure requirements in plain language, and the system would automatically translate the request into a complete provisioning plan.

Research is also needed to apply formal verification techniques to automatically generated infrastructure configurations. Such verification could ensure that deployments satisfy safety, reliability, and compliance requirements before they are executed. Another improvement involves developing multi-objective reinforcement learning systems that learn policy preferences directly from human operators, allowing the framework to adapt its decision-making priorities based on observed operational choices. Finally, the telemetry dataset collected during the six-month evaluation will be released as an open benchmark. This dataset will allow future research to compare different autonomous infrastructure management approaches under reproducible experimental conditions.

## REFERENCE

1. Roberts, T.A., Atwell, J., Sigler, E., & Doorn, Y.V. (2015). DevOps for VMware Administrators.
2. Dobaj, J., Krisper, M., & Macher, G. (2019). Towards Cyber-Physical Infrastructure as-a-Service (CPlaaS) in the Era of Industry 4.0. European Conference on Software Process Improvement.
3. Morris, K. (2016). Infrastructure as Code: Managing Servers in the Cloud.
4. Cleveland, S.B., Dooley, R., Perry, D., Stubbs, J., Fonner, J.M., & Jacobs, G.A. (2018). Building Science Gateway Infrastructure in the Middle of the Pacific and Beyond: Experiences using the Agave Deployer and Agave Platform to Build Science Gateways. Proceedings of the Practice and Experience on Advanced Research Computing: Seamless Creativity.
5. Soni, M. (2015). End to End Automation on Cloud with Build Pipeline: The Case for DevOps in Insurance Industry, Continuous Integration, Continuous Testing, and Continuous Delivery. 2015 IEEE International Conference on Cloud Computing in Emerging Markets (CECM), 85-89.
6. Wu, C.E., Yu, H., Jann, J., Burugula, R.S., Dubey, N., & Nguyen, M. (2017). Automation of Cloud Node Installation for Testing and Scalable Provisioning. Proceedings of the 10th IEEE/ACM International Conference on Utility and Cloud Computing Companion.

7. Sandobalín, J., Insfrán, E., & Abrahão, S.M. (2018). Towards Model-Driven Infrastructure Provisioning for Multiple Clouds. Integrated Spatial Databases.
8. Sandobalín, J., Insfrán, E., & Abrahão, S.M. (2017). End-to-End Automation in Cloud Infrastructure Provisioning. Integrated Spatial Databases.