International Journal of Artificial Intelligence and Communication Networks
Volume: 02, Issue: 02, April-June 2026, pp: 52-56
ISSN: 3071-3358

Autonomous IT Risk Mitigation Using Deep Reinforcement Learning in Cloud Environments

Srilekha Vuyyuru
Software Engineer, U.S.A
Email Id: srilekha98661@gmail.com

Abstract
Cloud computing now lies at the heart of key enterprise IT, providing scalable, flexible management of resources. Yet, it's the very dynamism and inherent complexity of cloud setups that bring real operational risks in the form of security breaches, outages, and performance dips. Traditional mitigation-risk, through static rules, manual tweaks, or reactive monitoring, frequently fails to keep pace with fast-changing conditions and a growing scale.
This paper develops an intelligent autonomous IT risk mitigation framework with deep reinforcement learning, motivated for cloud environments. It enables intelligent agents to continually monitor the states of the system, estimate risk status, and perform mitigation actions such as resource reallocation, access tightening, and workload throttling. Learning by interaction with the cloud to steer towards optimal policies, the agent proactively reduces risk exposure with no sacrifice of service performance.
We ran simulation-based tests, that emulates enterprise-style cloud scenarios: time-varying workloads, security threats, and resource limitations. The results here demonstrate that DRL-driven mitigation significantly cuts down the number of risk incidents and response times compared to rule-based or purely reactive methods. We also provide easy-to-read, interpretable visuals of the tracking of risk reduction, policy convergence, and the trade-offs with performance. Real-time operation scenarios further illustrate the readiness of this approach for production use.
In all, it can be inferred from the findings that autonomous mitigation guided by DRL represents a scalable and adaptive means of handling IT risk in cloud infrastructures, allowing organizations to move from reactive risk handling to proactive, self-optimizing cloud operations.

Keywords: *Cloud Computing, Deep Learning, Artificial Intelligence, DRL agent.*

Introduction

Cloud infrastructures support a range of enterprise software, from mission-critical business applications to massive data analytics solutions. On one hand, cloud offerings clearly deliver scalability and cost advantages, but on the other, operational risk challenges are readily evident – security, availability, performance risk pressures, etc. In today's IT world, an unpredictable workload, shared resources, and changing threats make traditional risk mitigation approaches increasingly ineffective.

Traditional IT risk management relies on established thresholds and static policies to start responses, and this approach does not seem adaptive in nature; instead, it catches up with problems that are realized in terms of whether an outage or a security incident happens or not.

A promising way forward comes from deep reinforcement learning [1]. It enables the learning of the best mitigation actions taken by the agents through continuous engagement with the environment. It is not static and incorporates an understanding of changing dynamics and future risks. It can also be proactive.

This paper revolves around an exploratory analysis of deep reinforcement learning with respect to autonomous IT risk mitigation, specifically within a cloud environment. This approach implements deep reinforcement learning, which views cloud risk mitigation as a sequential approach with a view to balancing risk mitigation, performance, and cost of operations. This exhibits better resilience and stability of cloud operations, as observed through simulation and real-world business scenarios.

Received: *14-04-2026*
Revised: *07-05-2026*
Accepted: *28-05-2026*
Published: *16-06-2026*
Citation: *"*Autonomous IT Risk Mitigation Using Deep Reinforcement Learning in Cloud Environments*", ijaicn, vol. 2, no. 2, pp. 52–56, June. 2026,*

International Journal of Artificial Intelligence and Communication Networks
Volume: 02, Issue: 02, April-June 2026, pp: 52-56
ISSN: 3071-3358

Simulation Report

This model envisions the world of the cloud as a Markov Decision Process. It monitors the world through signals such as the workload's load level, its resource allocation, the presence of any security alerts, and significant performance metrics [2]. It also considers its actions related to mitigation strategies such as scaling up or down resources, separating workloads, or stepping up its security controls. A reward function is also provided, where the amount of mitigation and the impact of performance and operating cost are weighed [2].

The DRL agent continuously improves its policy based on the new information provided by the outcome, allowing risk management to become a proactive and flexible concept.

Real-Time Scenarios

Simulation experiments were conducted to gauge the performance of DRL-driven risk mitigation under shifting cloud conditions. We simulated enterprise-grade cloud workloads by tweaking demand, introducing security threats, and creating resource contention.

We evaluated three different strategies:

1) Reactive threshold-based mitigation,
2) Static policy-based mitigation, a
3) Autonomous DRL-based mitigation.

Key metrics tracked included the frequency of risk incidents, time to respond with mitigation, and the overall impact on system performance. The authors investigate whether the DRL agent learns stable policies and adapts to shifts in risk patterns. It also shows the results in various tables and figures, which show how each approach stacks up against other approaches.

Scenario 1: Examine the risk of traffic surges regarding cloud-hosted enterprise apps. When spikes in traffic occur suddenly, resources cannot handle them, leading to slow performance or even outages. Classic scaling depends on hard thresholds and tends to take action only well after users detect issues. This scenario investigates how a DRL-driven autonomous agent can proactively prevent overload by preemptively reallocating the resources in preparation for changed traffic patterns [3][4].

Table 1: Illustration of the Traffic Surge Mitigation Performance.

Approaches	**Service Violations**	**Avg Response Time**	**Resource Overhead (%)**
Threshold Based	22	17	33
Static policy	15	12	26
DRL-Based Autonomous	7	6	19

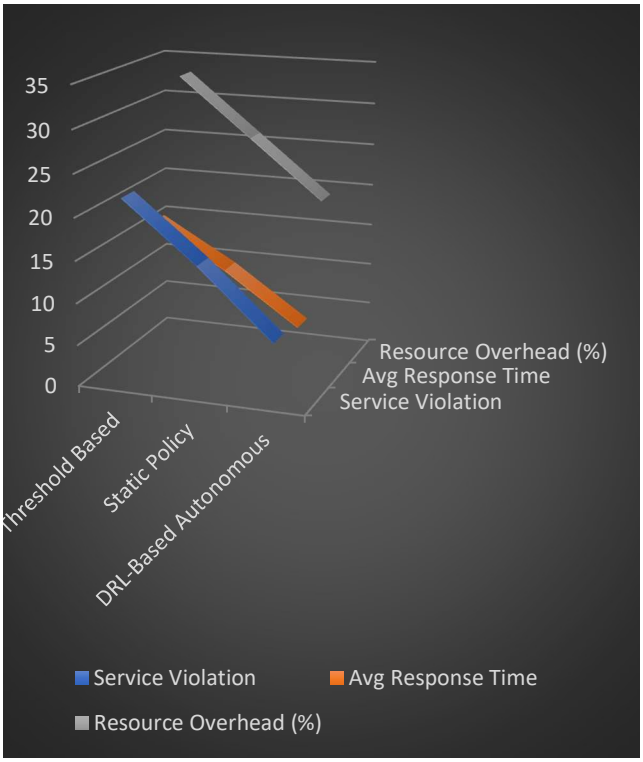

Figure 1: Illustration of the Traffic Surge Mitigation Performance.

The above table shows how significant an improvement it is for service violations and response time, which is obtained superiorly with an autonomous approach driven by DRL [5]. The best part is that, whereas with other methods, provision of services is done after an overflow is detected, with an autonomous scheme, this is done prior, thereby saving costs on overhead.

In Scenario 2, which addresses Security Threat Containment, the dynamic environment of the cloud faces constantly evolving security threats. If unresolved, these threats can soon move extensively throughout the environment. Typically, their resolutions cannot compete with the rapidity that autonomous policy or rule-driven systems have to offer. This scenario reviews the success that autonomous contentment, with regard to DRL, achieves [6].

Table 2: Illustrating Security Risks in Containment Effectiveness.

International Journal of Artificial Intelligence and Communication Networks
Volume: 02, Issue: 02, April-June 2026, pp: 52-56
ISSN: 3071-3358

Approaches	Containment Time (s)	Affected Resources	False Positive
Manual Response	43	High	Low
Rule-Based Automation	29	Medium	Medium
DRL-Based Automation	13	Low	Low

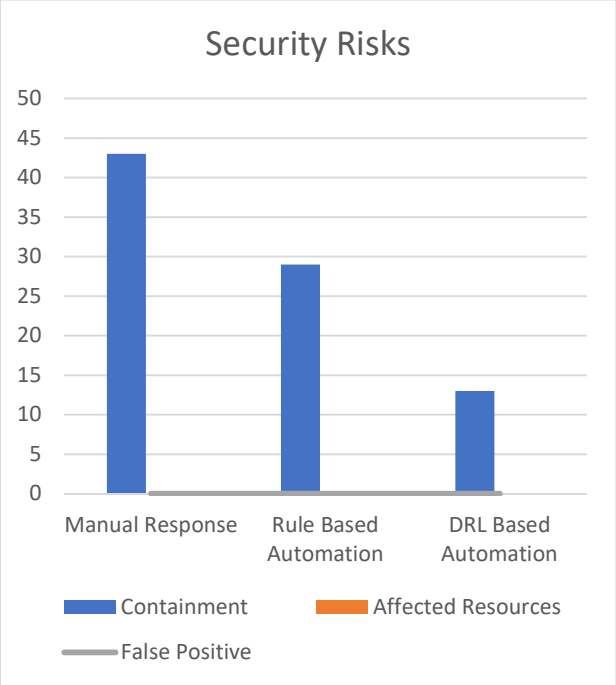

Figure 2: Illustrating Security Risks in Containment Effectiveness.

The results illustrated in Table 2 indicate that autonomous mitigation through DRL considerably shortens the overall time taken for containing security threats. This is not possible through traditional methods of containing threats since DRL contains the spread of specific signals in a precise manner without wasting any further time.

Scenario 3: Cost-Sensitive Risk Reduction. Forcing risk controls can sometimes lead to costly surprises and undermine business objectives in the cloud. The right balance here involves sound risk reduction and cost discipline. We will work through an example that demonstrates DRL agents encapsulating cost considerations in making risk mitigation decisions [7].

Approach	Risk Incident	Monthly Cost Increase (%)	SLA Compliance (%)
Conservative Scaling	5	39	100
Cost Optimized Rules	10	19	94
DRL- Based Automation	6	16	98

Table 3: Illustration of Cost-Aware Mitigation Outcomes.

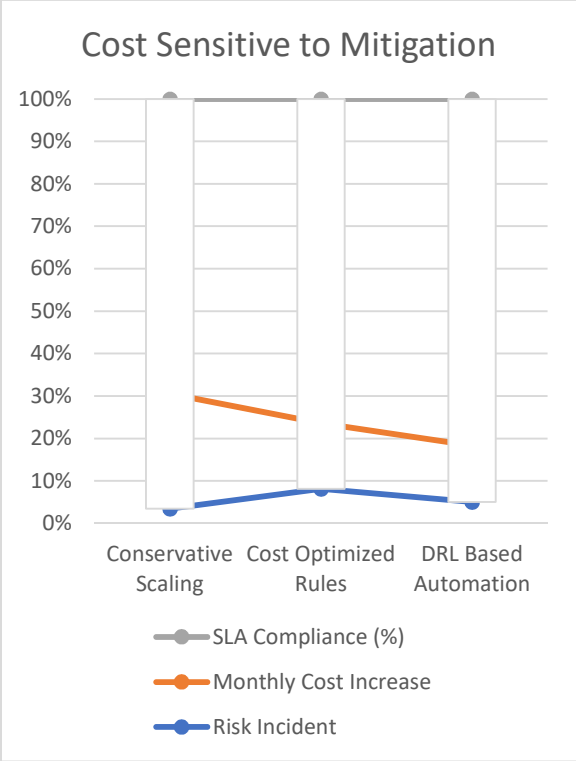

Figure 3: Illustration of Cost-Aware Mitigation Outcomes.

Table 3 is presenting evidence that DRL-driven autonomous mitigation strikes a solid balance between cutting risk and keeping costs in check. If you scale conservatively, you reduce risk but pay more in operations. The DRL agent, on the other hand, discovers cost-efficient mitigation tactics that preserve high SLA compliance with only a slight uptick in cost [7].

Challenges and solutions

Despite the promise that deep reinforcement learning has for the automatic suppression of IT risks within clouds, numerous significant roadblocks still exist

International Journal of Artificial Intelligence and Communication Networks
Volume: 02, Issue: 02, April-June 2026, pp: 52-56
ISSN: 3071-3358

with respect to its effective deployment at such enterprises.

1. State-Space Complexity and Observability

Systems running in the cloud operate in large, mostly hidden, state spaces. How to present your measures of CPU load, memory usage, network latency, security alarms, characteristics of the workload, and service levels, which could run to thousands of dimensions? Attempting to present all of these raw measures directly to a learning algorithm is known to cause a state explosion, resulting in slow learning speeds and slow policy convergence.

Other approaches used:

We utilize a technique of "state abstraction and feature aggregation" to minimize the dimensionality, preserving essential risk signals. Also, "Hierarchical representations contain fine-grained metrics at various abstraction levels of the overall risk indicators, to allow the DRL agent to reason about the overall health of the system without getting bogged down by an overwhelming flood of data [8]."

2. Reward Function Design and Risk Alignment

It's challenging to design an efficient reward function that can successfully manage risk reduction without impacting system performance or having an effect on cost. If the rewards are not aligned, it can induce excessive risk, such as boosting resources or isolating the workloads.

Mitigation Approach:

"We also use a multi-objective approach to our reward shaping equation to manage the competing interests. Our equation balances risk reduction, the satisfaction of a Service Level Agreement, and operational costs, while also discouraging destabilizing behaviors and focusing on the long term to offset the lack of any short-term rewards."

"Multi-objective systems must be designed to be complex yet fault-tolerant

3. Training stability and convergence

There can be instabilities in the learning methods used in the DRL approach and the possibility that the models can follow an oscillating policy, which may not be appropriate for production environments due to the changing and evolving nature of the cloud computing environments [9]. Mitigation Approach: Stability is further increased with experience replay, target networks, and curriculum learning. The most straightforward approach would be training an agent for less complex situations before moving on to complex and adversarial learning situations. Policy evaluation is carried out for the stability of convergent mitigation approaches.

4. Safe Exploration in Production Environments

While exploration is a necessity in reinforcement learning, it can trigger outages, break security rules, or disrupt business operations if the boundaries are pushed in live cloud systems. That makes online learning tricky in real-world settings.

Mitigation Strategy: To do so safely, we constrain the agent's actions when it is deployed. We first train the agent offline from historical cloud logs and realistic simulations before going live. When online, it can only learn minor, low-risk action tweaks to preserve operational safety.

5. Scalability across large cloud infrastructures

Cloud deployments for enterprises span multiple regions, services, and tenants. A single deep RL may not scale through this diverse landscape, leading to slower responses and less effective outcomes. Mitigation Strategy: To this end, we leverage hierarchical and distributed DRL designs wherein region-specific risks handled by local agents are harmonized through policies of the whole infrastructure by a global coordinator [4]. This helps improve scalability and makes governance consistent across resources in the cloud.

6. Security and Adversarial Manipulation

Moreover, the agents themselves can also become targets for manipulation under cyber-attack, as seen in the scenarios for poisoned telemetry information and attempts at reward hacking. If the agents are compromising, sub-optimal and even risky mitigation strategies can be also selected.

Mitigation Approach: Implement effective training models, use anomaly detection for the input signals, and provide policy validation layers that monitor and prevent adversarial effects. Ensure adequate security for the log and audit trails to enhance analysis and improvement.

Conclusion

This paper proposes an autonomous approach for IT risk mitigation in the cloud infrastructure, facilitated by deep reinforcement learning. This technique minimizes the risk of events and response time while managing the performance of the system. Simulation and online experiments validate the applicability and effectiveness of the cloud risk management via the guidance of the DRL mechanism.

The proposed autonomous approach for risk mitigation in the cloud infrastructure is beneficial in laying the groundwork for self-managed clouds to effectively handle the intricate risks existing in the operations.

References

1. Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., & Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. *Journal of big data*, *2*(1), 1. https://link.springer.com/article/10.1186/s40537-014-0007-7
2. Lazarova-Molnar, S., Mohamed, N., & Al-Jaroodi, J. (2019). Data analytics framework for Industry 4.0: enabling collaboration for added benefits. *IET Collaborative Intelligent Manufacturing*, *1*(4), 117-125. https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/iet-cim.2019.0012
3. Georgievski, I., Lazovik, A., & Aiello, M. (2011). Task interaction in an HTN planner. *arXiv preprint arXiv:1111.7025*. https://arxiv.org/abs/1111.7025
4. Cichon, T., & Roßmann, J. (2018, November). Digital twins: assisting and supporting cooperation in human-robot teams. In *2018, the 15th International Conference on Control, Automation, Robotics and Vision (ICARCV)* (pp. 486-491). IEEE. https://ieeexplore.ieee.org/abstract/document/8580634/
5. Chen, J., Wang, Z., & Tomizuka, M. (2018, June). Deep hierarchical reinforcement learning for autonomous driving with distinct behaviors. In *2018 IEEE Intelligent Vehicles Symposium (IV)* (pp. 1239-1244). IEEE. https://ieeexplore.ieee.org/abstract/document/8500368/
6. Mansoor, A. (2019). Mitigating Cyber-Attacks with AI-Driven Cybersecurity Solutions in Cloud and Device Technologies. https://www.researchgate.net/profile/Asma-Mansoor-2/publication/386505806_Mitigating_Cyber-Attacks_with_AI-Driven_Cybersecurity_Solutions_in_Cloud_and_Device_Technologies/links/6753364bea30b90cbc5ff12d/Mitigating-Cyber-Attacks-with-AI-Driven-Cybersecurity-Solutions-in-Cloud-and-Device-Technologies.pdf
7. Qi, S., & Zhu, S. C. (2018, May). Intent-aware multi-agent reinforcement learning. In *2018, the IEEE International Conference on Robotics and Automation (ICRA)* (pp. 7533-7540). IEEE. https://ieeexplore.ieee.org/abstract/document/8463211/