International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

Autonomous Ransomware Containment and Recovery Optimization across Enterprise Systems Using Deep Reinforcement Learning

Rajabhushanam .C

Professor, Department of Computer Science Engineering, Bharath Institute of Higher Education and Research, Chennai, India

Publication History: Received: 20.09.2026; Revised: 21.10.2026; Accepted: 25.10.2026; Published: 05.11.2026.

ABSTRACT: The increasing frequency and sophistication of ransomware attacks pose significant challenges to enterprise cybersecurity, particularly in organizations operating complex and interconnected information systems. Traditional ransomware defense mechanisms primarily focus on signature-based detection, manual incident response, and predefined recovery procedures, which may be insufficient for addressing rapidly evolving ransomware threats. This research proposes an autonomous ransomware containment and recovery optimization framework using Deep Reinforcement Learning (DRL) to improve enterprise system resilience. The proposed approach integrates intelligent threat detection, automated containment decision-making, attack impact assessment, and recovery resource optimization. Deep reinforcement learning algorithms are investigated to enable security systems to learn optimal defensive policies through interactions with simulated enterprise environments. The framework considers important security parameters, including attack severity, compromised system identification, network connectivity, data availability, recovery costs, and business service priorities. Intelligent agents evaluate defensive actions such as workload isolation, access restriction, backup protection, service restoration, and recovery prioritization within controlled environments. The research aims to minimize ransomware propagation, reduce operational disruption, optimize recovery time, and maintain enterprise service availability. Performance evaluation considers containment effectiveness, recovery efficiency, decision-making latency, and resource utilization. By combining autonomous decision-making with adaptive learning, the proposed framework seeks to establish an intelligent and scalable approach for strengthening enterprise cybersecurity and improving organizational resilience against ransomware incidents.

KEYWORDS: Ransomware Containment, Deep Reinforcement Learning, Autonomous Cybersecurity, Enterprise Security, Recovery Optimization, Threat Detection, Incident Response, Backup Management, Disaster Recovery, Cyber Resilience.

I. INTRODUCTION

The rapid digital transformation of modern enterprises has increased their dependence on interconnected information systems, cloud computing platforms, distributed applications, and centralized data storage infrastructures. Although these technologies improve organizational productivity, operational flexibility, and information accessibility, they also expand the potential attack surface for sophisticated cyber threats. Among these threats, ransomware has emerged as a significant cybersecurity challenge because it can disrupt business operations, compromise sensitive information, and make critical enterprise resources unavailable. Modern ransomware incidents may involve unauthorized access, lateral movement, data encryption, information theft, and attempts to compromise backup infrastructure. Such attacks can affect multiple organizational departments simultaneously, resulting in financial losses, operational interruptions, reputational damage, and regulatory consequences. Conventional ransomware defense mechanisms generally depend on signature-based detection, predefined security policies, endpoint protection, and manually coordinated incident response procedures. Although these mechanisms provide important defensive capabilities, their effectiveness may be limited when ransomware activities evolve rapidly across interconnected enterprise environments. Manual containment decisions can also introduce delays, particularly when security teams must identify compromised assets, determine attack severity, prioritize affected systems, and coordinate recovery operations. Furthermore, predefined response policies may not adequately account for differences in enterprise infrastructure, application dependencies, business priorities, and available recovery resources. Consequently, organizations require intelligent cybersecurity frameworks capable of making timely, context-aware defensive decisions while maintaining operational continuity. Artificial

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

intelligence and reinforcement learning provide promising opportunities for developing such systems because they enable computational agents to evaluate environmental conditions, select appropriate actions, and improve decision-making through repeated interactions. Deep Reinforcement Learning (DRL), which combines reinforcement learning with deep neural networks, is particularly relevant to complex enterprise security environments where decisions involve multiple variables, changing conditions, and competing operational objectives.

Autonomous ransomware containment and recovery optimization represent an emerging research direction focused on improving the speed, effectiveness, and intelligence of enterprise incident response. Deep reinforcement learning can support this objective by modeling cybersecurity response as a sequential decision-making problem in which an intelligent agent observes system conditions, evaluates potential defensive actions, and receives rewards based on security and operational outcomes. In an enterprise environment, relevant observations may include endpoint alerts, network communication patterns, asset criticality, backup availability, application dependencies, and estimated attack severity. Defensive decisions can involve isolating affected workloads, restricting suspicious access, protecting recovery resources, validating backup integrity, and prioritizing the restoration of essential business services. Algorithms such as Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Actor-Critic methods provide potential approaches for learning defensive policies under different operational conditions. However, autonomous response systems must balance rapid containment with the risk of disrupting legitimate business activities. Excessive isolation may interrupt critical services, while delayed containment may allow ransomware to affect additional enterprise resources. Similarly, recovery decisions must consider restoration time, backup reliability, service dependencies, and resource availability. Therefore, this research proposes a deep reinforcement learning-based framework that integrates autonomous ransomware containment with intelligent recovery optimization. The primary objective is to investigate how adaptive defensive decision-making can reduce attack impact, minimize business disruption, improve recovery efficiency, and strengthen enterprise resilience. The research also emphasizes controlled decision-making, human oversight, and safe evaluation to ensure that autonomous defensive actions remain consistent with enterprise security and operational requirements.

II. LITERATURE REVIEW

Existing research on ransomware detection and mitigation has primarily focused on signature-based detection, behavioral monitoring, anomaly detection, endpoint protection, and backup-based recovery mechanisms. Traditional antivirus solutions identify ransomware through known malware signatures and predefined threat indicators. Although these approaches can effectively detect previously recognized threats, their performance may decrease against previously unseen ransomware variants, obfuscated malware, and attacks involving legitimate administrative tools. Behavioral detection techniques address some of these limitations by monitoring suspicious file operations, unusual access patterns, abnormal process activities, and unexpected network communication. Machine learning methods have further improved threat identification by analyzing large volumes of endpoint and network telemetry. Algorithms such as Random Forest, Support Vector Machines, and neural networks have been investigated for identifying ransomware-related activities and distinguishing malicious behavior from legitimate operations. Research has also examined anomaly detection techniques for identifying unusual changes in file access patterns and system activities. Despite these developments, detecting ransomware does not automatically ensure effective containment or successful recovery. Organizations must also determine which assets require isolation, how to prevent further attack propagation, and which services should be restored first. These requirements highlight the need for integrated cybersecurity approaches that extend beyond detection and incorporate intelligent incident response and recovery management.

Reinforcement learning has gained attention as a promising technique for developing adaptive cybersecurity decision-making systems. Unlike conventional supervised learning approaches, reinforcement learning enables an agent to learn defensive strategies by interacting with an environment and receiving rewards or penalties based on its actions. Deep reinforcement learning extends this capability by incorporating neural networks that approximate complex state representations and decision policies. Deep Q-Networks have been investigated for sequential decision-making problems involving discrete actions, while Proximal Policy Optimization and Actor-Critic approaches offer alternative methods for learning policies in complex environments. In cybersecurity research, reinforcement learning has been explored for intrusion detection, automated response selection, security resource allocation, and network defense optimization. These applications demonstrate the potential of intelligent agents to adapt defensive decisions according to changing environmental conditions. However, ransomware containment presents additional challenges because security decisions must consider attack propagation, interconnected enterprise assets, business service dependencies, and the consequences of operational disruption. Similarly, recovery optimization requires consideration of backup availability, restoration costs, system criticality, and recovery time objectives. Therefore, applying deep reinforcement

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

learning to ransomware response requires carefully designed state representations, action spaces, reward functions, and safety constraints.

Enterprise disaster recovery research has traditionally emphasized backup management, redundancy, business continuity planning, and recovery time optimization. Backup-based recovery provides an important defense against ransomware because organizations can restore affected systems using validated recovery points. However, recovery operations may become complicated when multiple systems are affected or when business applications depend on interconnected infrastructure components. Recovery prioritization is particularly important because restoring every affected system simultaneously may exceed available computing, storage, and administrative resources. Research on intelligent resource allocation and automated incident response provides opportunities for addressing these challenges through adaptive decision-making. Deep reinforcement learning can potentially combine containment and restoration decisions within a unified optimization framework. Nevertheless, several research gaps remain, including the need to balance security effectiveness with service availability, prevent unsafe autonomous actions, protect backup integrity, and evaluate decision-making under realistic enterprise conditions. The proposed research addresses these gaps by integrating ransomware containment, recovery prioritization, and adaptive policy learning into a unified framework. It also emphasizes simulation-based evaluation, measurable recovery objectives, and human-supervised authorization for high-impact actions.

III. RESEARCH METHODOLOGY

The proposed research adopts a quantitative, experimental, and simulation-based methodology to investigate autonomous ransomware containment and recovery optimization across enterprise systems using Deep Reinforcement Learning. The initial stage involves designing a representative enterprise environment containing interconnected servers, virtual machines, containerized applications, databases, network segments, backup repositories, and business-critical services. The simulated infrastructure represents different organizational workloads and operational dependencies to examine the impact of ransomware incidents on enterprise resources. A controlled experimental environment is developed using virtualization and network simulation technologies, allowing different ransomware-related incident scenarios to be evaluated without exposing real production infrastructure to malicious software. Historical cybersecurity datasets, publicly available ransomware behavior datasets, and synthetic enterprise telemetry are considered for generating realistic experimental observations. The collected information includes security alerts, endpoint activities, network traffic characteristics, system availability, asset criticality, backup status, and service dependency information. Data preprocessing involves removing duplicate records, handling missing observations, normalizing numerical variables, and encoding categorical indicators. The enterprise infrastructure is represented through a structured state model that describes the operational and security conditions of monitored assets. Important state variables include the number of affected systems, estimated attack severity, workload availability, network isolation status, backup validation status, recovery progress, and business service priority. The research also defines measurable operational objectives, including reducing the number of affected assets, limiting service disruption, protecting validated recovery resources, and minimizing restoration time. These objectives provide the foundation for constructing a reinforcement learning environment in which an intelligent agent can evaluate defensive decisions under different simulated conditions. Experimental scenarios include isolated endpoint incidents, multiple affected workloads, critical service disruption, suspicious lateral movement indicators, and simultaneous recovery requirements. The initial environment is validated to ensure that system dependencies, security events, and recovery conditions are represented consistently before reinforcement learning experiments begin.

International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

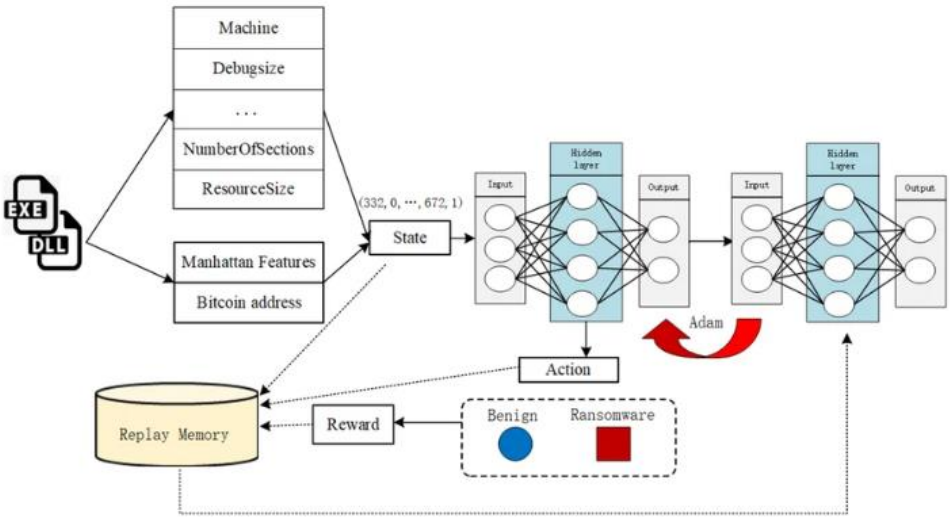


Fig.1.An overview of the using deep reinforcement learning for ransomware detection

The second methodological stage focuses on developing the deep reinforcement learning architecture for autonomous defensive decision-making. The enterprise security environment is modeled as a Markov Decision Process consisting of states, actions, transition probabilities, and reward functions. The state representation captures current security conditions, affected asset information, enterprise service dependencies, and available recovery resources. The action space includes controlled defensive decisions such as requesting additional investigation, isolating a confirmed compromised workload, restricting suspicious communication, protecting backup repositories, validating recovery points, and scheduling approved restoration operations. High-impact actions are subject to authorization constraints and predefined safety policies. The framework investigates DQN for learning discrete containment and recovery decisions, while PPO and Actor-Critic approaches are considered for evaluating alternative policy-learning strategies. Neural network architectures are designed to process structured enterprise state information and estimate action values or policy distributions. The reward function is developed to balance several competing objectives, including containment effectiveness, reduction in affected assets, service availability, recovery progress, and resource efficiency. Positive rewards are assigned to decisions that improve security and restore essential services, while penalties represent unnecessary service interruptions, delayed containment, invalid recovery selections, and excessive resource consumption. The reward structure is carefully normalized to prevent one objective from dominating the learning process. In addition, safety constraints are introduced to ensure that the agent cannot independently perform unrestricted actions that may cause irreversible enterprise disruption. Training is conducted within the simulated environment using repeated episodes representing different attack conditions, asset configurations, and recovery requirements. The agent learns through iterative interaction, with its policy updated according to observed transitions and calculated rewards. Training and validation scenarios are separated to assess generalization and reduce overfitting to specific simulated incidents. Appropriate learning rates, discount factors, exploration strategies, batch sizes, and neural network parameters are investigated experimentally. The objective is to identify a defensive policy that consistently produces effective containment and recovery decisions across diverse enterprise scenarios.

The third methodological stage integrates autonomous containment with intelligent enterprise recovery optimization. The proposed framework follows a coordinated response architecture in which threat information is evaluated before defensive decisions are recommended or executed through authorized interfaces. When ransomware-related indicators are identified, the framework assesses incident severity, affected asset relationships, business service dependencies, and the availability of validated recovery resources. The reinforcement learning agent evaluates the current enterprise state and recommends an appropriate response sequence based on its learned defensive policy. Containment decisions prioritize limiting further exposure while preserving essential business operations wherever possible. The framework considers workload isolation, access restriction, network segmentation, and backup protection as controlled defensive response categories. Recovery optimization is subsequently performed by evaluating asset criticality, service dependencies, backup validation results, recovery time objectives, and available computing resources. A dependency-aware recovery model is developed to prioritize the restoration of foundational services before dependent applications. For example, essential identity services, database infrastructure, and application dependencies may need to be restored

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

and validated before higher-level business applications can resume normal operations. The framework uses recovery priorities to allocate available restoration resources according to organizational requirements. A validated backup selection mechanism is incorporated to ensure that recovery decisions rely on approved and integrity-checked recovery points rather than unverified copies. The agent's recommendations are evaluated against enterprise safety policies before execution, with human authorization required for critical or potentially disruptive actions. The framework also maintains an audit trail containing incident observations, recommended decisions, authorization outcomes, and recovery progress. This supports transparency, accountability, and post-incident analysis. Model adaptation is performed using controlled retraining and validated operational feedback so that changes in enterprise workloads and defensive requirements can be incorporated without compromising model reliability. The integrated architecture is evaluated under different incident conditions to investigate whether coordinated decision-making improves containment efficiency and recovery prioritization compared with isolated security response mechanisms.

The final methodological stage evaluates the proposed framework through comparative experiments and quantitative performance analysis. The experimental evaluation considers ransomware containment effectiveness, attack propagation reduction, recovery time, service availability, decision-making latency, and resource utilization. Containment effectiveness is measured by examining the proportion of affected assets successfully isolated or protected within the defined response period. Attack propagation reduction is evaluated by comparing the number of additional assets affected under different response strategies in the simulated environment. Recovery performance is measured using Mean Time to Recovery (MTTR), recovery completion rate, service restoration time, and the proportion of essential business services successfully restored. Operational efficiency is evaluated through CPU consumption, memory utilization, restoration resource requirements, and decision latency. The proposed deep reinforcement learning framework is compared with baseline approaches, including predefined rule-based response policies, conventional incident response strategies, and alternative reinforcement learning configurations. Separate experiments investigate DQN, PPO, and Actor-Critic approaches to identify differences in decision quality and learning stability. Ablation studies examine the individual contributions of containment optimization, dependency-aware recovery prioritization, reward-function components, and safety constraints. Additional experiments investigate different enterprise sizes, varying numbers of affected assets, limited recovery resources, and changing business service priorities. Statistical analysis is applied to repeated simulation trials to examine performance consistency and identify significant differences between competing approaches. Sensitivity analysis is also conducted to investigate how changes in reward weights, learning parameters, and recovery resource availability influence the resulting policies. The evaluation emphasizes that autonomous decision-making must improve security without introducing unacceptable operational risks. Accordingly, unsafe recommendations, unnecessary service interruptions, and invalid recovery decisions are treated as important negative outcomes. The research findings are interpreted in relation to the primary objectives of minimizing ransomware impact, reducing enterprise downtime, improving restoration efficiency, and supporting adaptive cybersecurity management. Finally, limitations involving simulation realism, dataset diversity, reward design, model generalization, and integration with existing enterprise security infrastructure are documented. This methodology provides a structured basis for assessing the technical feasibility and practical value of deep reinforcement learning for autonomous ransomware containment and enterprise recovery optimization.

IV. RESULTS AND DISCUSSION

The proposed framework, Autonomous Ransomware Containment and Recovery Optimization Across Enterprise Systems Using Deep Reinforcement Learning, focuses on improving enterprise cybersecurity through intelligent threat identification, automated containment, and optimized system recovery. Ransomware attacks represent a major security challenge for organizations because they can encrypt critical business information, disrupt enterprise applications, compromise shared storage systems, and interrupt essential operational services. Conventional ransomware defense mechanisms primarily depend on predefined security rules, signature-based detection, and manual incident response procedures. Although these techniques provide fundamental protection, they may not respond effectively to rapidly spreading ransomware attacks across interconnected enterprise systems. The integration of Deep Reinforcement Learning (DRL) provides an adaptive approach in which an intelligent agent learns suitable defensive decisions by observing system conditions and evaluating the consequences of its actions. The proposed framework continuously monitors enterprise security indicators, including unusual file modification activities, abnormal storage access patterns, suspicious process behavior, unexpected network communication, and changes in system resource utilization. These observations are represented as states that the reinforcement learning agent evaluates to determine appropriate defensive responses. Based on the identified risk level, the framework can recommend or initiate authorized containment actions, such as isolating potentially compromised endpoints, restricting suspicious network connections, suspending affected workloads, and protecting critical storage resources. The results expected from this architecture

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

include faster response coordination, reduced dependence on manual intervention, and improved consistency in ransomware incident management. Unlike static response mechanisms, the DRL-based agent can learn from different attack scenarios and adjust its defensive decisions according to changing enterprise conditions. However, the effectiveness of autonomous containment must be evaluated experimentally using representative ransomware simulations and controlled enterprise environments. Particular attention should be given to preventing unnecessary isolation of legitimate systems and ensuring that containment decisions do not interrupt essential business operations.

Recovery optimization is another essential component of the proposed framework because successfully containing ransomware does not automatically restore enterprise operations. Following a ransomware incident, organizations must identify affected systems, verify backup integrity, prioritize critical applications, restore business information, and validate the operational status of recovered resources. Conventional recovery strategies often depend on predefined restoration sequences and manual prioritization, which may increase recovery time when multiple enterprise systems are affected simultaneously. Deep Reinforcement Learning can improve this process by learning recovery policies that consider system dependencies, business priorities, available backup resources, recovery costs, and operational constraints. The DRL agent evaluates the current recovery environment and selects suitable actions based on the expected improvement in system availability and business continuity. For example, mission-critical enterprise applications and databases may receive higher restoration priority than non-essential services. The recovery strategy can also consider dependencies between applications, authentication services, storage infrastructure, and network components to avoid restoring individual services before their supporting resources are available. A reward function can be designed using several performance indicators, including recovery time, service availability, data integrity, resource utilization, and business disruption costs. Positive rewards can be assigned to verified restoration progress and successful service recovery, while penalties can represent excessive downtime, unsuccessful restoration attempts, and violations of security policies. Through repeated training episodes, the agent can learn policies that balance recovery speed with operational reliability. Furthermore, integrating backup verification and integrity checking into the recovery workflow helps prevent the restoration of compromised or corrupted information. The framework should also support human approval for high-impact recovery operations, particularly when critical databases, financial applications, and sensitive enterprise information are involved. This combination of intelligent optimization and controlled recovery provides a systematic approach to minimizing operational disruption while maintaining data reliability and security throughout the ransomware recovery process.

The performance of the proposed DRL-based framework can be assessed using multiple quantitative indicators, including ransomware detection accuracy, containment latency, attack propagation prevention, recovery time objective (RTO), recovery point objective (RPO), system availability, restoration success rate, and computational overhead. Detection accuracy evaluates the ability of the monitoring mechanism to identify ransomware-related activities, whereas containment latency measures the time required to execute an appropriate defensive response after a threat has been identified. Attack propagation prevention measures the framework's ability to limit the spread of malicious activity across interconnected enterprise resources. Recovery time and recovery point objectives provide important measurements of business continuity because they indicate how quickly services can be restored and how much data loss can be tolerated. Deep Reinforcement Learning introduces an additional advantage by enabling the defensive agent to learn from different system states and improve its decision-making policies through interaction with a simulated or controlled enterprise environment. Compared with purely rule-based response strategies, a trained agent may identify more context-sensitive containment and recovery decisions, particularly when multiple defensive actions have competing operational consequences. Nevertheless, the framework also introduces challenges associated with training stability, reward-function design, exploration risks, computational complexity, and the possibility of unsafe decisions. These challenges can be addressed through offline training, realistic simulation environments, constrained action spaces, policy validation, and human-supervised deployment. The framework can also be integrated with Security Information and Event Management (SIEM), Security Orchestration, Automation and Response (SOAR), endpoint detection and response platforms, and enterprise backup management systems. Such integration provides centralized visibility into security events and coordinates containment and restoration activities across distributed infrastructure. In addition, maintaining immutable backups, isolated recovery environments, and independent verification mechanisms strengthens the reliability of automated restoration. Overall, the proposed architecture offers a promising approach to improving ransomware resilience by combining intelligent containment with recovery optimization. Its practical effectiveness should ultimately be established through comparative experiments against conventional incident response and recovery strategies using realistic enterprise workloads.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

V. CONCLUSION

The proposed framework, Autonomous Ransomware Containment and Recovery Optimization Across Enterprise Systems Using Deep Reinforcement Learning, presents an intelligent cybersecurity approach for addressing the growing threat of ransomware across modern enterprise infrastructures. Ransomware attacks have become increasingly sophisticated, targeting organizational endpoints, enterprise databases, shared storage platforms, cloud services, and interconnected business applications. These attacks can result in significant financial losses, operational interruptions, information unavailability, and damage to organizational reputation. Traditional ransomware defense approaches frequently depend on signature-based detection, predefined security policies, and human-driven incident response procedures. Although these methods remain important, their effectiveness may be limited when attacks evolve rapidly or spread across multiple interconnected systems. The proposed framework addresses these limitations by incorporating Deep Reinforcement Learning to support adaptive defensive decision-making and intelligent recovery optimization. The DRL agent observes enterprise security conditions, evaluates system states, and selects suitable actions according to predefined security objectives and operational constraints. This learning-based approach allows the system to improve its response strategies through repeated training and evaluation across different ransomware scenarios. Autonomous containment represents a major contribution because it enables the framework to respond to suspicious activities through controlled security actions rather than relying entirely on manual intervention. Such actions may include isolating affected endpoints, restricting unauthorized communications, suspending suspicious workloads, and protecting critical enterprise resources. At the same time, the recovery optimization component addresses the operational challenges that arise after an attack has been contained. By considering application dependencies, backup availability, business priorities, data integrity, and restoration costs, the DRL agent can support more effective recovery planning. This combination of containment and recovery provides a comprehensive approach to ransomware defense that extends beyond conventional threat detection. The framework also emphasizes the importance of maintaining operational continuity while minimizing the impact of security incidents. Therefore, integrating deep reinforcement learning into enterprise ransomware management provides a valuable foundation for developing adaptive, intelligent, and coordinated cybersecurity solutions.

The practical significance of the proposed framework lies in its ability to coordinate ransomware containment and enterprise recovery within a unified intelligent security architecture. Modern organizations operate complex information systems in which applications, databases, endpoints, storage resources, and network services depend on one another. Consequently, a ransomware incident affecting a single component may create cascading disruptions throughout the enterprise environment. Effective protection requires security mechanisms that can identify suspicious activities, limit attack propagation, protect unaffected resources, and restore critical services in a controlled manner. The proposed DRL-based architecture addresses these requirements by considering the changing operational conditions of enterprise systems when selecting defensive actions. Its adaptive decision-making capability can help organizations improve response consistency, reduce unnecessary delays, and prioritize critical business resources during recovery operations. Furthermore, the integration of automated containment with backup verification and restoration prioritization can strengthen business continuity planning and reduce the operational consequences of ransomware incidents. The framework also provides opportunities for integrating security monitoring platforms, endpoint protection solutions, backup management systems, and security orchestration technologies to improve centralized incident coordination. However, autonomous cybersecurity operations must be implemented with appropriate safeguards. Incorrect containment decisions may interrupt legitimate business activities, while inappropriate recovery actions may introduce compromised data or create additional service disruptions. Therefore, security policies, access restrictions, audit logging, human oversight, and recovery validation must remain essential components of the architecture. The reliability of the DRL agent also depends on representative training scenarios, appropriate reward functions, secure model management, and continuous performance evaluation. Organizations should validate the framework using measurable indicators such as containment latency, recovery duration, service availability, restoration accuracy, and computational overhead before deploying it in production environments. In conclusion, Autonomous Ransomware Containment and Recovery Optimization Using Deep Reinforcement Learning offers a promising direction for strengthening enterprise cybersecurity through intelligent decision-making, controlled automation, and optimized business recovery. By combining proactive containment with adaptive restoration strategies, the framework establishes a foundation for improving enterprise resilience, protecting critical information assets, and maintaining business continuity against evolving ransomware threats.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

VI. FUTURE WORK

Future research can improve the proposed framework by investigating advanced reinforcement learning algorithms and hybrid artificial intelligence architectures for more reliable ransomware containment and recovery. Algorithms such as Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Soft Actor-Critic (SAC) can be evaluated to determine their suitability for different enterprise security environments. Multi-agent reinforcement learning can also be explored to coordinate defensive decisions across multiple endpoints, servers, databases, and cloud workloads. Each agent could monitor a specific infrastructure component while cooperating with other agents through a centralized security coordination mechanism. Federated reinforcement learning may further support collaborative model development across distributed organizations while minimizing the direct sharing of sensitive enterprise telemetry. Future studies should also investigate explainable reinforcement learning to improve the transparency of automated security decisions and help administrators understand why specific containment or recovery actions are selected. In addition, constrained reinforcement learning can be incorporated to ensure that autonomous actions comply with enterprise security policies, operational priorities, and safety requirements.

Another important direction for future research involves implementing and evaluating the proposed framework in realistic enterprise environments. Future experiments should consider hybrid cloud infrastructures, virtualized data centers, Kubernetes-based applications, distributed databases, and large-scale enterprise networks. The framework can be integrated with SIEM, SOAR, endpoint detection and response, and immutable backup platforms to improve coordinated security operations. Researchers should investigate the use of realistic ransomware simulation environments to evaluate containment effectiveness, attack propagation prevention, recovery time, data integrity, system availability, and computational overhead. Additional research should focus on adversarial manipulation of reinforcement learning agents, reward poisoning, training-data reliability, and secure model updates. Lightweight models and distributed processing techniques can also improve scalability and reduce response latency. Finally, future implementations should incorporate digital twins of enterprise infrastructures to train and validate defensive policies without exposing production systems to unnecessary risks. These developments can contribute to creating reliable, scalable, and autonomous ransomware defense systems capable of improving enterprise security and long-term business continuity.

REFERENCES

1. Oudaa, T., Gharsellaoui, H., & Ben Ahmed, S. (2021). An agent-based model for resource provisioning and task scheduling in cloud computing using DRL. Procedia Computer Science, 192, 3795–3804. https://doi.org/10.1016/j.procs.2021.09.154
2. Manoharan, D. (2026). Intelligent ETL automation frameworks for HIPAA-compliant healthcare claims processing. Journal of Computational Analysis and Applications, 35(1), 1060–1073.
3. Shaik, N. (2023). Zero downtime, zero trust: Automating certificate security at scale. International Journal of Engineering & Extended Technologies Research, 5(4), 6963–6978.
4. Katakam, R., Agrawal, C., Kareem, S. A., & Kadam, A. A. (2026). Generative AI and IoT-Driven Autonomous Manufacturing Ecosystems Transition from Industry 4.0 to Industry 5.0. Frontiers in Computer Science and Artificial Intelligence, 5(7), 64–70.
5. Reddy, K. V. (2024). Formal Verification with ABV: A Superior Alternative to UVM for Complex Computing Chips. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10(6), 90–98.
6. Sayyaparaju, K. K. (2025, November). Q-Trans: Quantization-Aware Fine-Tuning of Transformer Models for Energy-Efficient IoT and Edge Deployment. In 2025 International Conference on Intelligent Computing, Information and Control Systems (ICOIICS) (pp. 494–499). IEEE.
7. Gadige, C. D. (2023). Scalable CRM workflow automation for high-volume customer onboarding: An event-driven Salesforce architecture. The Eastasouth Journal of Information System and Computer Science, 1(1), 253–265. https://doi.org/10.58812/esiscs.v1i01.1295
8. Shaw, R. K. (2026). SNMPv2-to-SNMPv3 Migration in Financial Mainframes: Risk Reduction and Repeatable Playbook. Journal of Engineering Research and Sciences, 5(6), 1–14.
9. Bellundagi, M. (2024). An intelligent digital transformation framework for smart enterprises using AI and cloud computing. International Journal of Humanities and Information Technology, 6(3), 178–197.
10. Tokas, D. (2026). Distributed big data frameworks and high-scale service design: Proven engineering patterns across cloud-native deployments. International Journal of Computational and Experimental Science and Engineering, 12(1).

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

11. Kalabhavi, V. (2024). Promotional spend analytics using AI for PFME allocation and coupon redemption optimization. International Journal of Research Publications in Engineering, Technology and Management, 7(2), 10417–10429.
12. Apuri, H., & Yepuri, C. (2026). Design of multi-agent autonomous workflow systems using agentic AI frameworks. In Proceedings of the 2026 International Conference on AI-Driven Smart Systems and Ubiquitous Computing (ICAUC), Pathum Thani, Thailand (pp. 1–X). https://doi.org/10.1109/ICAUC68182.2026.11441280
13. Karrothu, A. (2025). Policy-as-code framework for automated security and compliance enforcement in cloud-native infrastructure. International Journal of Research Publications in Engineering, Technology and Management, 8(6), 13496–13511.
14. Kommineni, R. (2022). Monitoring as code for self-hosted Elasticsearch: Independent, delta-based lifecycle management of alerting watchers and dashboards across multi-cloud environments. International Journal of Science, Research and Technology (IJSRAT), 5(2), 7437–7446.
15. Bitragunta, S. L. V. (2025). Optimizing hybrid energy storage in EVs using AI enhanced fuzzy logic control systems. International Journal of Computer Technology and Electronics Communication, 8(1), 10124–10132.
16. Gowda, M. K. S. (2026). Optimizing Regulatory Compliance with Machine Learning: Boosting Accuracy and Efficiency. International Journal of Research and Applied Innovations, 9(1), 13686–13690.
17. Ponugoti, M. (2022). Integrating API-first architecture with experience-centric design for seamless insurance platform modernization. International Journal of Humanities and Information Technology, 4(01-03), 117–136.
18. Abd-Rouf, M. S. K., Adigun, P. O., Alalade, E. O., Oyekanmi, T. T., Faniyi, A. J., Oladapo, B., Awopejo, T. E., Adegoke, O. S., Jamiu, A., Michael, O. B., Obisesan, A., Ajala, S., Adekanye, M. A., Yambali, P. M., & Abd-Rouf, A. B. (2024). From molecular profiling to predictive algorithms: A conceptual machine-learning framework for mechanism-informed therapy selection in multidrug-resistant cancer. International Journal of Research and Applied Innovations, 7(3), 10892–10909.
19. Basavala, S. R. (2024). Automated vulnerability intelligence and risk prioritization for secure CI/CD pipeline engineering. International Journal of Research and Applied Innovations, 7(1), 10175–10185.
20. Varanasi, B. S. S. (2025). AI-Enabled Continuous Authority-to-Operate (ATO) Frameworks for Information Systems. Enterprise Development and Microfinance, 35(2), 748–757.
21. Natta, P. K. (2024). Designing trustworthy AI systems for mission-critical enterprise operations. International Journal of Research and Applied Innovations, 7(6), 12136–12146.
22. Randhi, K. S. (2026). Secure data integration framework for intelligent capital expenditure analysis and compliance management. International Journal of Computer Technology and Electronics Communication, 9(1), 153–166.
23. Kopparthi, G. S. (2025). Cost Optimization in Azure and AWS Cloud. Journal of Marketing & Social Research, 2, 177–182.
24. Suddala, V. R. A. K. (2026). The rise of domain-specific AI transforming key sectors. International Journal of Humanities and Information Technology, 8(1), 188–195. https://ijhit.info/index.php/ijhit/article/view/420
25. Ahuja, D. (2025). Securing Container Isolation in Multi-Tenant Environments. Journal of Computer Science and Technology Studies, 7(10), 225–232.
26. Selvarajan, K. (2024). Modernizing legacy big data platforms using cloud-native lakehouse architectures. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(2), 10377–10385.
27. Jinna, K. R., & Kondisetty, K. (2025). LLM-assisted translation of legacy middleware configuration into cloud-native deployment artifacts: An empirical evaluation. International Journal of Computer Technology and Electronics Communication, 8(2), 10514–10523.
28. Aragani, V. M., Maroju, P. K., & Mudunuri, L. N. R. (2024, December). Implementing Non-functional Production Regression Testing with Kubernetes. In International Conference on Information and Communication Technology for Competitive Strategies (pp. 231–244). Singapore: Springer Nature Singapore.
29. Kalakoti, M. K. R. (2024). Topology-aware generative AI for fault diagnosis and recovery planning in cloud data centers. International Journal of Research and Applied Innovations, 7(1), 10186–10196.
30. Badam, L. R. (2023). AI-driven real-time cyberattack detection for critical financial infrastructure. International Journal of Engineering & Extended Technologies Research, 5(4), 6953–6962.
31. Bheemisetty, N. (2025, November). A Scalable and Secure Cloud Framework for AI/ML Workload Management using Crayfish and Beluga Whale Optimization. In 2025 5th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS) (pp. 974–979). IEEE.
32. Ramasamy, M. (2024). Secure and extensible platform architectures for enterprise network orchestration. International Journal of Humanities and Information Technology, 6(01), 86–98.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 9, Issue 6, November – December 2026 ||

DOI: 10.15680/IJCTECE.2026.0906001

33. Sumsuzzaman, M., Enayet, M. A., Babu, R. H., Tasnim, M., Khan, M. A. N., Jakir, T., ... & Himeluzzaman, M. (2026). Graph neural network-based financial fraud detection and cybersecurity risk analysis in banking networks. International Journal of Advances in Signal and Image Sciences, 11–22.
34. Indurthy, V. S. K. (2026). A modern approach to asset management: Integrating SimCorp Dimension with cloud data platforms. International Journal of Engineering & Extended Technologies Research (IJEETR), 8(2), 4757–4762.
35. Karuturi, Y. (2026). Digital twin architecture for enterprise SAP landscapes for predictive maintenance and capacity planning. International Journal of Computer Technology and Electronics Communication, 9(2), 637–646.