Computing, Information Systems, Development Informatics & Allied Research Journal
Vol. 17 No. 3, September 2026 - www.cisdijournal.net

A Conceptual and Design-Science Framework for Autonomous Self-Healing Endpoint Detection and Response Systems

[1]Danquah Paul & [2]Matey Akwetey Henry
[1]Council for Scientific and Industrial Research – Institute for Scientific and Technological Information (CSIR-INSTI), Ghana
[2]University of Professional Studies, Accra (UPSA), Ghana
E-mail: pdanquah@csir.org.gh; henraclesgh@yahoo.com
Phone Nos: +233243784082; +233244572485.

ABSTRACT

The rising complexity of cyber threats has shown the inadequacies of traditional Endpoint Detection and Response (EDR) systems, which depend significantly on static detection protocols and human involvement for incident investigation, reaction, and recovery. This paper introduces an Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR) Framework that integrates Multi-Agent Artificial Intelligence, Multi-Agent Deep Reinforcement Learning (MADRL), and autonomous recovery mechanisms to strengthen enterprise cyber resilience. The research employs a Design Science Research (DSR) methodology, underpinned by an extensive review of current literature, to create and develop an innovative cybersecurity artefact that includes specialised intelligent agents tasked with endpoint monitoring, threat detection, forensic investigation, decision-making, autonomous response, recovery, and knowledge coordination. The suggested paradigm conceptualises cyber security as a collaborative, decentralised, and perpetually adaptive ecosystem that adjusts to changing attack patterns while reducing human involvement. The framework aims to enhance detection precision, minimise the average time for detection and response, enable autonomous endpoint recovery, and promote ongoing organisational learning via reinforcement learning. The study presents a theoretically informed and practically applicable architectural framework that enhances autonomous cyber defence through the integration of adaptive intelligence, collaborative decision-making, and self-healing capabilities, thus establishing a basis for the development of next-generation intelligent EDR systems and future empirical validation in enterprise settings.

Keywords: Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR), Multi-Agent Artificial Intelligence, Multi-Agent Deep Reinforcement Learning (MADRL)

CISDI Journal Reference Format
Danquah Paul & Matey Aketey Henry (2026): A Conceptual and Design-Science Framework for Autonomous Self-Healing Endpoint Detection and Response Systems. Computing, Information Systems, Development Informatics & Allied Research Journal. Vol 17 No 3, Pp 41-62.. Available online at www.isteams.net/cisdijournal. dx.doi.org/10.22624/AIMS/CISDI/V17N3P3

1. INTRODUCTION

The increasing sophistication of cyberattacks has exposed limitations in conventional Endpoint Detection and Response (EDR) systems, particularly their dependence on predefined detection rules and human-led incident response. Emerging threats, including polymorphic malware, fileless attacks, living-off-the-land techniques, and AI-enhanced attacks, challenge static detection mechanisms and increase the operational burden on Security Operations Centres (SOCs).

Although contemporary EDR platforms support endpoint monitoring, threat detection, forensic investigation, and containment, reliance on human analysts for critical response decisions can delay mitigation and increase operational costs. Artificial intelligence (AI) offers opportunities to develop more adaptive and autonomous cyber-defence systems. Multi-Agent Systems (MAS) support distributed security operations through specialized agents that exchange contextual information and coordinate activities such as behavioural monitoring, threat analysis, investigation, and containment. Complementing this approach, Multi-Agent Reinforcement Learning (MARL) enables agents to learn defensive strategies through interaction with dynamic environments and feedback, reducing dependence on exclusively predefined security policies (Zhu et al., 2024; Ning & Xie, 2024).

Reinforcement learning further provides a basis for adapting defensive decisions to changing attack behaviours and operational conditions. However, integrating learning-driven defence into enterprise environments presents unresolved challenges involving scalability, explainability, agent coordination, and adversarial robustness. Recent studies have examined advances in MARL, federated learning, and autonomous cyber-defence architectures, while identifying these challenges as barriers to practical deployment (Luo & Jiang, 2025; Jing et al., 2025). Another emerging development is the application of agentic AI to cybersecurity, enabling intelligent software agents to reason, plan, and coordinate defensive actions with reduced human intervention. Such capabilities create opportunities to automate incident handling and support cyber resilience through faster detection, response, and recovery. Nevertheless, existing research frequently examines detection algorithms, reinforcement learning techniques, and multi-agent coordination as separate capabilities, leaving scope for a unified, endpoint-centric architecture that integrates these functions across the incident response lifecycle.

To address this gap, this study proposes an Autonomous Intelligent Agent-Based Endpoint Detection and Response Framework that integrates MAS, reinforcement learning, and an Autonomous Response Engine. The conceptual framework comprises specialized agents that collaboratively monitor endpoints, detect suspicious activities, investigate incidents, coordinate containment, and initiate recovery. Reinforcement learning supports the continuous refinement of defensive decisions using feedback from operational experiences. By integrating distributed intelligence, adaptive decision-making, and autonomous response, the proposed framework seeks to establish a foundation for more resilient endpoint protection with minimal human intervention.

This research aims to present a framework that assists researchers in addressing the following research question:

How can a multi-agent AI architecture be designed to support autonomous detection, investigation, containment, recovery, and learning across the EDR lifecycle?

Addressing this question contributes to the advancement of autonomous cyber defence by presenting an adaptive, collaborative, and self-learning EDR architecture capable of improving detection accuracy, reducing response time, minimizing analyst workload, and strengthening organizational cyber resilience against emerging and previously unseen threats.

2. LITERATURE REVIEW

2.1 Evolution of Endpoint Detection and Response

Endpoint Detection and Response (EDR) has become fundamental to contemporary enterprise cybersecurity due to the escalating complexity of cyber-attacks that increasingly bypass conventional signature-based antivirus systems. Traditional endpoint protection solutions (EPPs) predominantly emphasised malware prevention using static signatures and heuristic analysis. The swift increase of advanced persistent threats (APTs), ransomware, fileless malware, and living-off-the-land (LotL) attacks has revealed the shortcomings of prevention-focused security frameworks. These emerging threats utilise legitimate system tools and previously unidentified weaknesses, making conventional antivirus technology insufficient for safeguarding modern enterprise environments (Kaur et al., 2024).

To rectify these deficiencies, EDR transformed from a reactive malware detection instrument into an all-encompassing endpoint security platform proficient in continuous monitoring, behavioural analysis, forensic investigation, threat hunting, and incident response. Contemporary EDR systems perpetually gather endpoint telemetry, encompassing process execution, registry alterations, network connections, authentication occurrences, memory activities, and file system interactions. The data are examined with behavioural analytics, machine learning, and cyber threat intelligence to identify malicious actions that may not conform to previously recognised attack signatures and methods (Kaur et al., 2024).

The incorporation of artificial intelligence (AI) and machine learning (ML) has markedly enhanced the functionalities of EDR solutions. Instead of depending exclusively on established indicators of compromise (IoCs), AI-driven EDR platforms can detect aberrant behaviours, correlate intricate attack sequences, and facilitate predictive threat detection. Mohamed (2025) contends that AI and ML have revolutionised cybersecurity through adaptive threat detection, automated incident analysis, and ongoing learning from evolving assault patterns, thus enhancing the efficacy of contemporary cyber defensive systems. Nonetheless, contemporary AI-driven EDR solutions predominantly function as decision-support instruments that necessitate human analysts to verify alarms, do forensic investigations, and sanction containment and recovery measures.

A notable advancement in the development of EDR is its incorporation into comprehensive detection and response frameworks, such as Security Information and Event Management (SIEM), Security Orchestration, Automation and Response (SOAR), and Extended Detection and Response (XDR) platforms. The integrations have enhanced cross-domain visibility by correlating telemetry from endpoints, networks, clouds, identities, and applications, facilitating more thorough threat detection and coordinated responses. Nonetheless, despite these advancements, endpoint security continues to rely significantly on established playbooks and the skill of analysts, especially during investigation, decision-making, and post-incident recovery (Shen et al., 2024).

Recent investigations have underscored the intrinsic limits of current EDR technology. Despite advancements in behavioural analytics and machine learning enhancing detection efficacy, commercial EDR platforms persist in producing substantial quantities of security warnings necessitating manual validation, hence exacerbating analyst fatigue and escalating operating expenses.

Moreover, advanced adversaries are increasingly utilizing evasion techniques tailored to circumvent endpoint monitoring, such as indirect system calls, memory-only execution, and hook manipulation, indicating that contemporary EDR architectures continue to be susceptible to sophisticated attack methodologies (Carvalho Junior, 2024). As a result, the progression of EDR is transitioning from intelligent detection to autonomous cyber defence. Recent studies support the integration of autonomous AI agents, reinforcement learning, and adaptive decision-making systems that can autonomously identify, analyse, mitigate, and recover from cyber incidents with limited human involvement. These features signify the forthcoming evolutionary phase of endpoint protection, evolving EDR from an analyst-dependent solution into an autonomous, self-learning, and self-healing cyber resilience platform. Despite increasing interest in autonomous cyber defence, existing EDR systems have not yet completely incorporated distributed intelligence, collaborative autonomous agents, continuous learning, and automatic recovery into a cohesive framework. This constraint serves as the principal impetus for the Autonomous Self-Healing Endpoint Detection and Response Framework suggested in this research.

2.2 Multi-Agent Systems in Autonomous Cyber Defence: Multi-Agent Systems for Autonomous Cyber Defence

The growing complexity and decentralised characteristics of enterprise information systems have expedited the implementation of Multi-Agent Systems (MAS) as a viable framework for autonomous cyber protection. A Multi-Agent System consists of several autonomous, intelligent agents that sense their surroundings, make independent decisions, interact with each other, and collaborate to attain common security goals. In contrast to centralised security architectures, multi-agent systems (MAS) allocate computational intelligence across various specialised agents, enhancing scalability, fault tolerance, adaptability, and resilience against advanced cyber attacks (Vyas et al., 2025). Contemporary enterprise environments comprising cloud infrastructures, Internet of Things (IoT) ecosystems, edge computing platforms, and hybrid networks are fundamentally decentralised, rendering Multi-Agent Systems (MAS) especially adept at safeguarding diverse endpoint ecosystems via coordinated and distributed decision-making.

In cybersecurity, MAS has been utilised for intrusion detection, malware analysis, cyber deception, vulnerability management, attack simulation, and autonomous network protection. Individual agents are generally allocated specific duties, including endpoint monitoring, anomaly detection, forensic investigation, threat intelligence correlation, policy enforcement, containment, and recovery. By means of ongoing communication and collaboration, these agents jointly attain security objectives that would be challenging to achieve with monolithic or centralised security solutions. The decentralised structure of MAS eradicates single points of failure and facilitates swift local reactions, while preserving global situational awareness via knowledge dissemination (Vyas et al., 2025).

Recent advancements in artificial intelligence have augmented MAS by including autonomous reasoning, collaborative planning, and adaptive learning functionalities. The advent of agentic AI has allowed intelligent software agents to autonomously detect cyber threats, analyse attack developments, orchestrate defensive measures, and implement reaction tactics with minimal human involvement. Recent research demonstrates that multi-agent architectures enhance operational flexibility by enabling specialised agents to dynamically assign tasks, share contextual information, and collaboratively optimise defensive strategies in rapidly changing adversarial settings (Zhang et al., 2025).

Notwithstanding these advancements, the implementation of MAS in enterprise cyber protection continues to be disjointed. Current research has primarily concentrated on network intrusion detection, cyber ranges, autonomous blue teaming, software-defined networking, and cyber-physical systems, however there is a relative paucity of studies examining endpoint-centric autonomous defence. Vyas et al. (2025) noted that while autonomous cyber-defence agents exhibit substantial potential, numerous obstacles persist concerning coordination mechanisms, continual learning, explainability, governance, and implementation in practical organisational settings. Recent evaluations of autonomous agent architectures similarly cite interoperability, trust management, communication efficiency, and policy compliance as significant barriers to practical adoption.

A significant drawback pertains to the breadth of current MAS-based cybersecurity systems. Most architectures emphasise discrete defensive capabilities, such as intrusion detection or automated response, rather than facilitating the entire incident response lifecycle. As a result, capabilities such as autonomous forensic investigation, coordinated endpoint isolation, intelligent recovery, and ongoing post-incident learning are often lacking. This functional fragmentation restricts the capacity of existing MAS architectures to provide comprehensive autonomous cyber resilience, especially in enterprise endpoint contexts where attacks frequently progress swiftly across various devices and stages.

These observations indicate a substantial research opportunity. Although MAS offers a robust basis for distributed cyber defence, there is a paucity of research on integrated architectures that merge collaborative intelligent agents with adaptive learning and autonomous recovery in Endpoint Detection and Response (EDR) systems. This study proposes a Autonomous Self-Healing Endpoint Detection and Response Framework that enhances current MAS research by incorporating specialised intelligent agents, collaborative decision-making, continuous learning via Multi-Agent Deep Reinforcement Learning, and autonomous recovery into a cohesive cyber-defence architecture that supports the entire endpoint incident response lifecycle.

.

2.3 Reinforcement Learning for Adaptive Cyber Defence: Reinforcement Learning and Multi-Agent Reinforcement Learning in Cybersecurity

Reinforcement Learning (RL) has emerged as a very promising paradigm in artificial intelligence for cybersecurity, as it allows autonomous agents to acquire optimal defensive methods through ongoing interaction with dynamic and adversarial environments. In contrast to supervised learning, which relies on labelled datasets and fixed decision boundaries, reinforcement learning utilises a trial-and-error approach wherein an agent observes its environment, executes an action, receives a reward or penalty, and modifies its policy to optimise long-term cumulative rewards. The adaptive learning feature renders reinforcement learning especially appropriate for cybersecurity, since attack methodologies change swiftly and novel threats often render static detection models obsolete (Abouhawwash, 2024).

Recent research indicate that reinforcement learning has been effectively utilised in intrusion detection, malware classification, vulnerability management, moving target defence, cyber deception, attack-path optimisation, and autonomous incident response. Through ongoing engagement with simulated or operational environments, reinforcement learning agents can ascertain optimal defence strategies without the necessity for extensive rule engineering or

annotated assault datasets. This capability signifies a substantial improvement over traditional machine learning methods, which frequently necessitate retraining when adversaries show new attack behaviours (Finistrella et al., 2025).

Notwithstanding these advancements, conventional single-agent reinforcement learning demonstrates significant limits in the context of enterprise cybersecurity. Contemporary cyberattacks generally transpire across decentralised infrastructures that include endpoints, cloud services, mobile devices, Internet of Things (IoT) devices, and hybrid networks. A solitary learning agent frequently lacks adequate environmental awareness and computing scalability to synchronise defensive measures across diverse contexts. Moreover, cybersecurity issues are intrinsically multi-faceted and partially discernible, necessitating concurrent monitoring, analysis, containment, and recovery actions that surpass the limitations of singular decision-making frameworks (Landolt et al., 2025).

The aforementioned restrictions have fostered an increasing interest in Multi-Agent Reinforcement Learning (MARL), which enhances reinforcement learning by allowing several autonomous agents to collaborate in the development of coordinated defence strategies. In MARL environments, agents perceive local conditions, perform autonomous behaviours, share contextual knowledge, and collaboratively enhance common security goals. This decentralised learning model closely resembles enterprise cyber-defence operations, as numerous security components function simultaneously across distant infrastructures. Finistrella et al. (2025) categorise MARL applications based on deployment environment and learning strategy, determining that collaborative reinforcement learning provides substantial benefits in scalability, adaptability, and resilience compared to traditional single-agent methods. Specifically, MARL facilitates distributed defence strategies that may coordinate intrusion detection, attack mitigation, and containment across several network domains while preserving local autonomy.

Recent methodological advancements have enhanced the applicability of MARL to autonomous cyber defence. Hierarchical MARL architectures partition intricate defence tasks into specialised sub-policies, enabling individual agents to concentrate on functions such as endpoint investigation, network monitoring, host recovery, or threat containment, while a superior policy orchestrates the overarching defensive strategy. Experimental assessments in authentic cyber-range settings indicate that hierarchical Multi-Agent Reinforcement Learning (MARL) enhances convergence velocity, defensive efficacy, and policy generalisation in complex adversarial scenarios marked by partial observability and deceptive attack strategies (Singh et al., 2025).

Likewise, research on autonomous cyber-defence indicates that reinforcement learning agents can proficiently safeguard enterprise networks from advanced persistent threats (APTs) by perpetually adjusting reaction strategies to evolving attacker behaviour. Instead of depending on manually created playbooks, RL-based defenders acquire containment techniques through interactions in simulated corporate environments, enhancing reaction consistency and diminishing reliance on human operators (Foley et al., 2024). Despite these encouraging advancements, numerous significant challenges persist. Contemporary research in Reinforcement Learning (RL) and Multi-Agent Reinforcement Learning (MARL) primarily emphasises network-centric applications, such as intrusion detection systems, software-defined networking, autonomous penetration testing, and attack mitigation.

Relatively scant focus has been allocated to endpoint-centric autonomous defence, wherein intelligent agents are required to synchronise behavioural monitoring, forensic analysis, endpoint isolation, remediation, and post-incident recovery. Furthermore, numerous published research depend on synthetic benchmark environments that insufficiently represent the complexity, diversity, and operational limitations of enterprise endpoint ecosystems (Finistrella et al., 2025; Mavroudis et al., 2025). Further research problems encompass computational complexity, communication overhead among collaborating agents, adversarial robustness, explainability, incentive engineering, and policy transfer across diverse organisational environments. Recent methodological guidance underscores persistent challenges with secure exploration, governance, trustworthiness, and the incorporation of RL-based decision-making into practical cybersecurity processes (Mavroudis et al., 2025).

Consequently, while reinforcement learning (RL) and multi-agent reinforcement learning (MARL) have laid a robust groundwork for adaptive cyber defence, current research has not yet presented a comprehensive framework that integrates collaborative intelligent agents, reinforcement learning, autonomous investigation, response orchestration, self-healing recovery, and continuous organisational learning within an Endpoint Detection and Response (EDR) architecture. This study's primary contribution is the proposal of a Autonomous Self-Healing Endpoint Detection and Response Framework that integrates Multi-Agent Artificial Intelligence with Multi-Agent Deep Reinforcement Learning to facilitate adaptive, collaborative, and resilient endpoint protection with minimal human involvement.

2.4 Research Gap

Autonomous Response Mechanisms and Self-Reparative Cyber Defence Although advancements in artificial intelligence have significantly enhanced cyber threat identification, reaction automation remains a relatively underdeveloped component of enterprise cybersecurity. Traditional Endpoint Detection and Response (EDR) and Security Orchestration, Automation, and Response (SOAR) systems predominantly depend on established playbooks and rule-based procedures to implement containment and remediation measures. While these methodologies diminish manual labour for recognised attack scenarios, they demonstrate restricted adaptability when faced with zero-day exploits, polymorphic malware, advanced persistent threats (APTs), and AI-driven attacks. Recent research has consequently focused on autonomous cyber defence systems that can independently reason, plan, and execute defensive operations in dynamic
and adversarial contexts (Vyas et al., 2025).

Autonomous Response Engines (AREs) signify the next advancement in cyber protection by transcending static orchestration to facilitate intelligent decision-making through automation. Instead of following predefined workflows, AREs utilise artificial intelligence, reinforcement learning, and agent-based reasoning to assess environmental conditions, anticipate attack developments, choose optimal response strategies, and continuously refine their policies based on operational feedback. This adaptive feature allows response mechanisms to develop in tandem with adversarial conduct, hence decreasing mean time to detect (MTTD), mean time to reply (MTTR), and analyst burden. Recent systematic evaluations identified autonomous reasoning, continuous learning, explainability, and realistic training settings as essential prerequisites for functional autonomous cyber-defence systems (Vyas et al., 2025).

The idea of self-healing cyber defence is closely associated with autonomous response, wherein security systems autonomously return compromised assets to their trusted operating states after successful containment. Self-healing systems encompass more than just threat detection; they integrate integrity verification, configuration restoration, vulnerability mitigation, automatic patch deployment, rollback to trusted configurations, and ongoing post-recovery monitoring. Agentic artificial intelligence has expedited this transition by facilitating autonomous software agents that can perceive, reason, deploy defensive tools, and orchestrate intricate recovery operations with minimal human involvement. Recent surveys demonstrate that agentic AI markedly improves cyber resilience through autonomous adaptation and ongoing operational learning inside distributed infrastructures.

Notwithstanding these advancements, the actual application of autonomous response engines is still constrained. Current research primarily focuses on autonomous detection, network defence, or orchestration in isolation, with minimal emphasis on the integration of autonomous detection, investigation, response, recovery, and organisational learning into a cohesive endpoint protection framework. Moreover, significant issues persist concerning decision explainability, governance, safety, policy compliance, and evaluation in authentic company settings. These constraints indicate that autonomous response ought to be seen not as a standalone capability but as a component of a cohesive self-healing cyber-defence framework that perpetually learns from operational experiences and evolves its defensive techniques over time.

2,.5 Critical Research Gap and Conceptual Synthesis

The examined literature indicates significant advancements in the utilisation of artificial intelligence, Multi-Agent Systems (MAS), reinforcement learning (RL), and autonomous cyber defence. Modern EDR platforms have transitioned from signature-based malware detection to behavior-driven security solutions, while MAS has enhanced distributed decision-making and RL has facilitated adaptive policy optimisation in dynamic threat settings. Nonetheless, these advancements are disjointed, as the majority of research focuses on isolated elements of autonomous cyber security rather than providing a cohesive and operationally feasible framework.

The initial research gap pertains to the restricted autonomy of existing EDR systems. Despite advancements in machine learning for behavioural detection and threat prioritisation, commercial platforms still depend significantly on human analysts for alarm validation, forensic investigation, containment decisions, remediation, and recovery. As a result, the entire incident response lifecycle is only partially automated, constraining the scalability and responsiveness of enterprise Security Operations Centres (SOCs) amidst the escalating volume and complexity of cyber attacks (Vyas et al., 2025).

A secondary gap pertains to the implementation of Multi-Agent Reinforcement Learning (MARL). Current MARL research has predominantly concentrated on intrusion detection, autonomous penetration testing, cyber ranges, software-defined networking, and network defence. Relatively few research examine endpoint-centric architectures in which numerous intelligent agents collectively engage in continuous monitoring, forensic investigation, adaptive reaction, autonomous recovery, and organisational learning. The potential of MARL to enable the entire Endpoint Detection and Response lifecycle remains inadequately investigated.

Third, contemporary autonomous cyber-defence infrastructures hardly integrate self-healing capabilities as a fundamental element of cyber resilience. Most frameworks conclude upon threat containment, with recovery operations either excluded or assigned to external operational procedures. Recent study underscores that forthcoming autonomous cyber-defence systems must incorporate recovery, continuous learning, explainability, and governance to ensure reliable and robust autonomous functionality.

Ultimately, substantial methodological obstacles persist. Current research often assesses reinforcement learning agents in simplified simulation models that fail to accurately reflect the intricacies of business endpoint ecosystems. Challenges such as communication efficiency among autonomous agents, adversarial robustness, policy transferability, explainable decision-making, ethical governance, and safe deployment continue to impede operational adoption.

This study proposes a Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR) Framework that conceptually integrates the advantages of Multi-Agent Artificial Intelligence, Multi-Agent Deep Reinforcement Learning, autonomous response mechanisms, and self-healing cyber resilience into a cohesive architecture. The suggested system integrates endpoint monitoring, collaborative threat detection, autonomous investigation, adaptive response, intelligent recovery, and ongoing organisational learning into a unified cyber-defence lifecycle, in contrast to existing models that highlight separate capabilities. This synthesis offers a theoretical advancement in autonomous cybersecurity research and a realistic architectural basis for next-generation intelligent endpoint protection solutions.

2.6 Synthesis of Reviewed Literature

Table 1: Synthesis Table: Existing Cybersecurity Approaches versus ASH-EDR

Approach	Existing capabilities and limitations reported in the literature	ASH-EDR capability or proposed advancement
EDR	Provides continuous endpoint monitoring, behavioural analytics, threat hunting, forensic investigation, and response. However, it remains dependent on human analysts for alert validation, investigation, and response approval.	Integrates autonomous investigation, adaptive response, self-healing recovery, and continuous learning into the endpoint protection lifecycle.
SOAR	Automates incident response through predefined playbooks and rules. Its adaptability is limited when facing unfamiliar or evolving attack scenarios.	Proposes adaptive response decisions informed by collaborative agents and Multi-Agent Deep Reinforcement Learning (MADRL), rather than relying exclusively on static playbooks.
XDR	Correlates telemetry across endpoints, networks, cloud environments, identities, and applications. Nevertheless, response workflows and post-incident recovery may remain dependent on established procedures and human expertise.	Extends coordinated security operations towards autonomous endpoint investigation, containment, remediation, and recovery.

Approach	Existing capabilities and limitations reported in the literature	ASH-EDR capability or proposed advancement
Multi-Agent Systems (MAS)	Distributes intelligence among specialised agents, supporting collaboration, scalability, fault tolerance, and local responses. Existing research often addresses isolated defensive functions rather than the complete endpoint incident lifecycle.	Coordinates specialised monitoring, detection, investigation, response, and recovery agents within a unified endpoint-centric architecture.
Reinforcement Learning (RL)	Enables adaptive defensive decisions through interaction, feedback, and reward optimisation. Single-agent RL may struggle with decentralised environments, partial observability, and coordinated defence across multiple endpoints.	Incorporates adaptive learning into a distributed defence architecture to support context-sensitive response decisions.
Multi-Agent Reinforcement Learning (MARL)	Supports collaborative learning and coordinated defence across distributed environments. Existing studies predominantly focus on network defence, intrusion detection, cyber ranges, and related applications; operational endpoint recovery remains insufficiently integrated.	Proposes collaborative adaptive decision-making across endpoint security agents, extending MARL towards the entire EDR lifecycle.
Autonomous Response Engines (AREs)	Use intelligent reasoning, planning, and automated actions to improve response speed and reduce analyst workload. However, autonomous detection, response, and recovery are frequently studied separately, with continuing concerns about governance and explainability.	Integrates autonomous investigation, response orchestration, recovery, and post-incident learning into a cohesive self-healing framework.
ASH-EDR (proposed)	Addresses the identified fragmentation among intelligent detection, distributed decision-making, adaptive learning, autonomous response, and recovery.	Proposes a unified framework combining specialised agents, MADRL, autonomous investigation and response, self-healing recovery, and continuous organisational learning, with minimal human involvement.

The principal contribution of ASH-EDR is not that its individual components are entirely new, but that it proposes to integrate capabilities commonly examined separately into a unified, endpoint-centric, adaptive, and self-healing architecture. These are proposed architectural capabilities; their effectiveness requires empirical validation against existing approaches.

3. METHODOLOGY

3.1 Research Approach and Conceptual Framework Development

This study adopts a conceptual framework development approach informed by the principles of Design Science Research (DSR). DSR provides a suitable methodological foundation because the study seeks to conceptualise an innovative cybersecurity artifact that addresses a complex organizational problem. Unlike empirical research approaches primarily concerned with observing or explaining existing phenomena, DSR emphasizes the development of purposeful artifacts, including models, methods, frameworks, architectures, and software systems, to address practical problems while generating scientific knowledge (Hevner et al., 2004; Peffers et al., 2007). In this study, the artifact is the proposed Autonomous Self-Healing Endpoint Detection and Response Framework utilizing Multi-Agent Artificial Intelligence (ASH-EDR). The conceptual framework addresses the limitations of existing Endpoint Detection and Response (EDR) systems, particularly their dependence on human analysts, predefined response procedures, fragmented autonomous capabilities, and limited integration of adaptive decision-making with automated recovery. The research therefore seeks to establish a coherent architectural model that integrates Multi-Agent Systems (MAS), Multi-Agent Deep Reinforcement Learning (MADRL), autonomous response mechanisms, and self-healing recovery within a unified endpoint-centric cyber-defence framework.

The selection of a conceptual approach is consistent with the design-oriented nature of DSR, which supports the formulation of artifacts and the development of generalizable design knowledge (Hevner et al., 2004; Peffers et al., 2007). Contemporary DSR emphasizes the importance of iterative reasoning, methodological transparency, and the alignment of artifact design with relevant organizational problems (Hevner et al., 2024). Furthermore, recent methodological extensions highlight the importance of incorporating intelligent and adaptive capabilities while maintaining traceability, governance, and explainability in AI-enabled artifacts (Tuunanen et al., 2024; Verma et al., 2026). These principles inform the conceptualization of ASH-EDR as a distributed, adaptive, and self-healing endpoint security architecture.

3.2 Conceptual Framework Development Process

The development of ASH-EDR is structured around four interrelated activities: problem characterization, derivation of conceptual requirements, architectural synthesis, and formulation of design propositions. These activities are informed by the design-oriented principles of Hevner et al. (2004) and the systematic artifact-development logic described by Peffers et al. (2007). However, the process is adapted to the objectives of a conceptual framework paper rather than a study involving prototype implementation or empirical performance evaluation.

(a) Problem Characterization and Motivation

The first activity establishes the cybersecurity problem addressed by the proposed framework. Drawing on the literature reviewed in this study, the analysis identifies limitations in existing EDR, Security Orchestration, Automation and Response (SOAR), Extended Detection and Response (XDR), MAS, reinforcement learning (RL), and Multi-Agent Reinforcement Learning (MARL) approaches. Particular attention is given to fragmented incident response capabilities, reliance on human intervention, limited integration of adaptive decision-making, and insufficient incorporation of autonomous recovery into endpoint security architectures. This problem characterization establishes the conceptual justification for an integrated self-healing EDR framework.

(b) Derivation of Conceptual Requirements

The second activity translates the identified research gaps into functional and non-functional requirements for ASH-EDR. The functional requirements encompass continuous endpoint monitoring, collaborative threat detection, autonomous investigation, adaptive response selection, coordinated containment, recovery, and post-incident learning. The non-functional requirements include scalability, resilience, explainability, secure agent coordination, policy compliance, and operational accountability. These requirements provide the design criteria against which the proposed architecture and its constituent components are conceptually organized.

(c) Architectural Synthesis and Framework Formulation

The third activity integrates the identified capabilities into a unified conceptual architecture. ASH-EDR is organized around specialized intelligent agents responsible for endpoint monitoring, threat detection, investigation, adaptive response, and recovery. The agents are conceptually coordinated through collaborative information exchange and a knowledge-sharing mechanism, while MADRL provides the basis for adaptive decision-making and the refinement of defensive policies through feedback. Autonomous response mechanisms translate selected defensive decisions into containment and remediation actions, while the self-healing component supports restoration to trusted operational states and subsequent monitoring. The architecture further incorporates governance, explainability, and policy enforcement as cross-cutting considerations intended to constrain autonomous actions and support accountable operation. This synthesis distinguishes ASH-EDR from approaches that concentrate on individual defensive capabilities by proposing an integrated conceptual model covering the endpoint incident response lifecycle. The resulting framework represents a theoretically grounded architectural proposition rather than a validated operational system.

(d) Formulation of Design Propositions and Research Contributions

The final activity articulates the conceptual relationships among the framework's components and identifies the expected contributions of their integration. The principal proposition is that coordinating specialized security agents with collaborative adaptive learning, autonomous response, and self-healing recovery can provide a more comprehensive architectural basis for autonomous endpoint defence than approaches that address these capabilities separately. Additional propositions concern the potential contribution of distributed decision-making to scalability, adaptive learning to responsiveness against evolving threats, and automated recovery to cyber resilience. These propositions are conceptual expectations derived from the framework's design rationale and require subsequent empirical investigation before their effectiveness can be established. The resulting contribution comprises the ASH-EDR conceptual architecture, its functional and non-functional requirements, the relationships among its principal components, and a set of design propositions to guide future implementation and evaluation. Consistent with design-oriented research, the framework is intended to contribute to the development of generalizable design knowledge while addressing a practical cybersecurity problem (Hevner et al., 2004; Peffers et al., 2007).

3.3 Methodological Considerations and Limitations

Although the conceptual approach provides a structured basis for formulating ASH-EDR, several methodological considerations remain. First, the proposed architecture and its anticipated benefits are derived from conceptual synthesis and design reasoning rather than empirical validation.

Consequently, claims concerning detection effectiveness, response speed, scalability, resilience, and recovery performance should be interpreted as propositions requiring future verification rather than demonstrated outcomes. Second, the integration of autonomous agents and MADRL raises questions concerning coordination overhead, secure information exchange, policy conflicts, reward design, and the reliability of adaptive decisions. Explainable artificial intelligence, ethical AI governance, adversarial robustness, privacy preservation, and compatibility with heterogeneous enterprise environments also require further examination. Third, autonomous containment and recovery introduce operational risks when defensive decisions are incorrect or when restoration mechanisms interact with compromised systems. Future research should therefore investigate mechanisms for policy enforcement, action authorization, recovery verification, and appropriate human oversight. These considerations are integral to the conceptual framework because autonomous operation must be balanced with accountability, operational safety, and organizational security requirements.

In summary, the methodology establishes a design-oriented conceptual foundation for ASH-EDR by connecting identified cybersecurity research gaps with explicit requirements, an integrated multi-agent architecture, and theoretically grounded design propositions. The resulting framework contributes a structured basis for future prototyping, empirical validation, and refinement without presupposing that the proposed capabilities have already been implemented or experimentally demonstrated.

3.4 Basis and Justification for the Proposed System Architecture

The proposed Autonomous Intelligent Agent-Based Endpoint Detection and Response (AIA-EDR) architecture aims to rectify various deficiencies in current Endpoint Detection and Response (EDR) systems. Despite substantial advancements in endpoint visibility via behavioural monitoring and machine learning, most commercial EDR platforms still rely predominantly on centralised analysis, manually crafted detection rules, and human analysts for incident investigation, response orchestration, and system recovery. As enterprise environments get more distributed and cyber threats evolve, human-dependent procedures lead to extended incident response times, alert fatigue, and variable response quality (Finistrella et al., 2025; Vyas et al., 2025).

The suggested architecture is based on the principle of autonomous cyber defence, wherein intelligent software agents continuously monitor the environment, analyse emerging threats, coordinate defensive measures, and enhance their performance through learning. This paradigm signifies a notable shift from traditional EDR solutions that primarily support security analysts instead of substituting monotonous operational tasks. Recent systematic reviews suggest that forthcoming cyber-defence platforms must integrate autonomous agents with capabilities for observation, reasoning, collaboration, and adaptive reaction to effectively address more sophisticated cyber-attacks (Vyas et al., 2025).

The core design choice was the implementation of a Multi-Agent System (MAS) architecture. In contrast to centralised cybersecurity platforms, MAS disaggregates intricate security operations into multiple specialised autonomous agents that collaborate to attain a unified security goal. This methodology provides numerous benefits, such as enhanced scalability, fault tolerance, parallel processing, and resilience to single points of failure. Enterprise infrastructures generally consist of thousands of diverse endpoints located across on-premises networks, cloud platforms, mobile devices, and Internet of Things (IoT) environments.

A decentralised design facilitates security decisions at the protected endpoint while permitting global knowledge exchange among collaborating agents. Recent surveys on Multi-Agent Reinforcement Learning (MARL) highlight decentralised decision-making as a key advantage of autonomous cyber-defence systems, enabling multiple agents to coordinate responses to distributed attacks while ensuring operational scalability (Finistrella et al., 2025; Landolt et al., 2025).

The Endpoint Monitoring Agent (EMA) serves as the foundational layer of the proposed architecture. Ongoing telemetry capture is crucial as contemporary attacks increasingly leverage legal administration tools, fileless malware, credential theft, and living-off-the-land tactics that never produce conventional malware signatures. Thus, endpoint visibility should encompass not just static file examination but also process execution, registry alterations, authentication occurrences, memory usage, user activity, and network interactions. Ongoing behavioural monitoring supplies the contextual knowledge necessary for astute threat analysis and facilitates the adaptive learning mechanisms utilised by advanced agents (Finistrella et al., 2025).

The Threat Detection Agent (TDA) conducts behavioural anomaly detection through the integration of machine learning and cyber threat intelligence. This component rectifies a significant limitation of rule-based EDR solutions, specifically their failure to identify novel attack patterns without human rule modifications. Reinforcement learning and behavioural analytics empower threat detection models to adjust to evolving adversarial tactics instead of solely depending on past signatures. Finistrella et al. (2025) assert that reinforcement learning offers a potential approach for improving detection, mitigation, and response in dynamic cybersecurity contexts, as defensive strategies may adapt continuously to changing threat situations.

Upon detection, the Investigation Agent (IA) conducts independent forensic analysis. Current EDR solutions often produce alerts but delegate incident reconstruction, attack attribution, and evidence correlation to human analysts. Isolating the inquiry of an autonomous intelligent agent enhances both modularity and explicability. The Investigation Agent reconstructs attack timelines, correlates indicators of compromise, detects potential attack pathways, predicts lateral movement, and supplies contextual evidence that underpins autonomous respoense judgements. This nuanced thinking corresponds with the development of autonomous cyber-defence frameworks that promote the use of specialised intelligent agents rather than singular security
systems (Loevenich et al., 2025).

The core intelligence of the proposed architecture is the Reinforcement Learning Decision Agent (RLDA). Reinforcement learning was chosen because of the highly dynamic, stochastic, and adversarial nature of cybersecurity situations. In contrast to supervised learning algorithms, reinforcement learning enhances its performance by ongoing contact with its environment by optimising cumulative rewards linked to effective defensive behaviours. This functionality allows the proposed system to autonomously adjust to new attack strategies without necessitating comprehensive retraining upon the emergence of new malware strains. Additionally, the design employs Multi-Agent Reinforcement Learning (MARL), enabling multiple intelligent agents to collaborate while concurrently acquiring coordinated defence techniques. Recent studies indicate that Multi-Agent Reinforcement Learning (MARL) markedly enhances scalability, collaboration, and resilience in autonomous cyber-defence settings relative to single-agent methodologies (Kiely et al., 2025; Morris et al., 2025).

The Autonomous Response Agent (ARA) converts acquired policies into implementable defensive measures. Conventional Security Orchestration, Automation, and Response (SOAR) platforms often implement established playbooks developed manually by cybersecurity specialists. While successful against known attack scenarios, static playbooks struggle to adapt to polymorphic malware, zero-day exploits, or AI-assisted attacks. The suggested response engine utilises policy-driven decision-making based on reinforcement learning, allowing for the dynamic selection of response measures based on attack severity, endpoint context, corporate security policies, and projected future outcomes. Recent experimental studies have shown that reinforcement learning agents can effectively coordinate containment strategies in realistic enterprise cyber-defence environments, while decreasing response latency and enhancing overall defensive performance (Morris et al., 2025; Kiely et al., 2025).

In contrast to most current autonomous defence frameworks, the proposed architecture includes a specialised Recovery Agent (RA). Recovery is a crucial, but somewhat underexamined, stage of incident response, as it embodies the primary goal of cyber resilience. Effective threat containment does not ensure the restoration of endpoint integrity or the resumption of organisational operations. Thus, the Recovery Agent conducts integrity verification, configuration restoration, vulnerability repair, rollback to trusted system states, and ongoing post-recovery monitoring. Incorporating recovery into the autonomous decision cycle converts the framework from a traditional detection-and-response system into a self-healing cybersecurity architecture that can restore operational continuity after effective containment (Loevenich et al., 2025).

The ultimate architectural element is the Knowledge Coordination Layer, facilitating effective communication among all intelligent agents. Instead of continuously exchanging raw telemetry, agents disseminate contextual threat intelligence, forensic evidence, established policies, reaction outcomes, and environmental insights. This cooperative knowledge-sharing system enhances decision-making quality while minimising communication burdens. Effective inter-agent communication has been recognised as a critical factor affecting the efficacy of MARL-based cybersecurity systems, as cooperative learning relies on prompt information exchange without overburdening distributed enterprise infrastructures (Zhu et al., 2024; Landolt et al., 2025).

The proposed architecture consolidates distributed intelligent agents, collaborative reinforcement learning, autonomous response orchestration, and self-healing recovery into a cohesive EDR framework. This architecture integrates complementary technologies to rectify the primary deficiencies noted in contemporary autonomous cyber-defence literature, establishing a scalable framework for adaptive endpoint protection that functions with minimal human oversight (Finistrella et al., 2025; Vyas et al., 2025).

4. PROPOSED FRAMEWORK FOR AUTONOMOUS SELF-HEALING ENDPOINT DETECTION AND RESPONSE UTILIZING MULTI-AGENT ARTIFICIAL INTELLIGENCE

The proposed framework introduces a decentralized, intelligent, and self-healing Endpoint Detection and Response (EDR) architecture that leverages collaborative Multi-Agent Artificial Intelligence (MAAI) to autonomously detect, investigate, contain, recover, and continuously learn from cybersecurity incidents. Unlike conventional EDR solutions that primarily support security analysts, the proposed framework enables intelligent software agents to cooperate throughout the complete cyber-defence lifecycle with minimal human intervention.

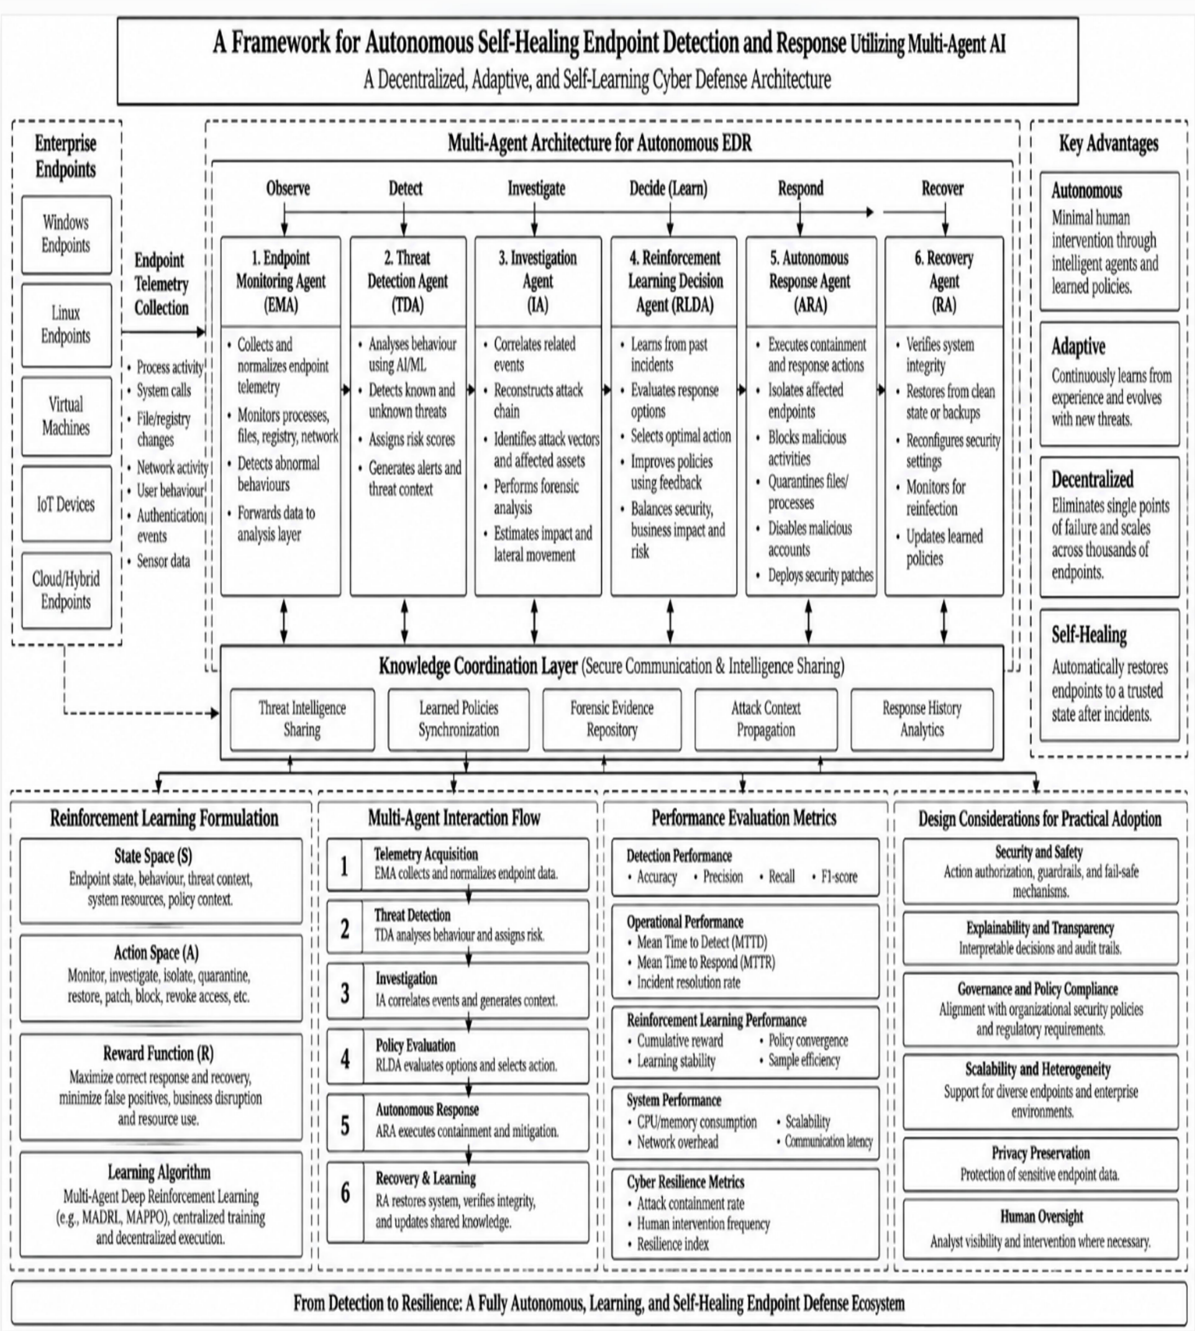


Figure 1: Framework for Autonomous Self-Healing Endpoint Detection and Response Utilizing Multi-Agent Artificial Intelligence
Source: Author(s)

The framework is founded upon five complementary design principles:

1. Distributed Intelligence – Security decisions are decentralized across specialized autonomous agents rather than relying on a single decision engine.
2. Collaborative Multi-Agent Coordination – Specialized agents exchange contextual knowledge to improve collective decision-making.
3. Adaptive Reinforcement Learning – Defensive policies evolve continuously through Multi-Agent Deep Reinforcement Learning (MADRL).
4. Autonomous Self-Healing – The framework automatically restores compromised endpoints to trusted operational states after successful containment.
5. Continuous Organizational Learning – Knowledge acquired from every security incident is retained to improve future defensive performance.

Collectively, these principles establish an intelligent cyber-defence ecosystem capable of adapting to evolving cyber threats while minimizing analyst workload and response latency.

4.1 Layered Framework Architecture

The proposed framework is organized into six logical layers that collectively realize the autonomous self-healing cyber-defence lifecycle.

Layer 1 – Endpoint Intelligence Layer

The Endpoint Intelligence Layer serves as the observational basis for the proposed Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR) framework. The main duty is to consistently oversee business endpoints and provide thorough situational awareness of their operational conditions. Each secured endpoint contains a Endpoint Monitoring Agent (EMA) that persistently collects and preprocesses security telemetry produced from process execution, system calls, memory usage, registry alterations, user activities, authentication events, network interactions, and file system operations. Instead of sending unprocessed telemetry directly to other agents, the EMA standardises, refines, and converts the gathered data into organised security observations appropriate for advanced analysis. This localised preprocessing reduces communication cost while maintaining the contextual information necessary for precise behavioural analysis, thus facilitating scalable and effective monitoring in distributed organisational situations.

Layer 2 - Cognitive Threat Assessment Layer

The Intelligent Threat Analysis Layer converts normalised endpoint telemetry into actionable security intelligence via the cooperative functioning of two specialised intelligent agents. The Threat Detection Agent (TDA) utilises behavioural analytics, machine learning, and cyber threat intelligence to detect anomalous behaviours, identify suspicious behavioural patterns, calculate dynamic threat scores, and assess attack severity. Upon the identification of suspicious behaviour, the Investigation Agent (IA) commences independent forensic analysis by reconstructing attack timelines, identifying attack vectors, correlating Indicators of Compromise (IoCs), assessing the probability of lateral movement, evaluating the overall severity of the incident, and producing contextual forensic evidence to facilitate subsequent response decisions. Distinguishing behavioural detection from forensic inquiry improves both clarity and decision precision by guaranteeing that autonomous response choices are underpinned by thorough contextual analysis instead of just anomaly detection.

Layer 3 – Autonomous Decision Intelligence Layer

The Autonomous Decision Intelligence Layer functions as the cognitive nucleus of the proposed system, generating intelligent defensive decisions through adaptive learning instead of relying on predefined security protocols. At the core of this layer is the Reinforcement Learning Decision Agent (RLDA), which perpetually assesses the current security posture of each endpoint by analyzing endpoint status, historical attack data, organisational security policies, contextual threat intelligence, environmental factors, and projected long-term rewards. The decision-making process is defined as a Markov Decision Process (MDP), with each state reflecting the current security status of an endpoint, determined by factors including CPU utilisation, active processes, registry changes, network activity, authentication events, file integrity, behavioural anomaly scores, threat intelligence indicators, and historical response data.

According to these observations, the RLDA independently determines the most suitable defensive measures, which encompass ongoing surveillance, evidence acquisition, alert issuance, termination of malicious processes, quarantine of executables, isolation of endpoints, blocking of malicious IPs, suspension of accounts, deployment of security patches, system restoration, or release of endpoints post-verification. The reinforcement learning policy aims to enhance long-term cyber resilience by incentivising precise threat detection, successful containment, prompt response, efficient recovery, and minimal false-positive rates, while imposing penalties for undetected attacks, erroneous containment decisions, superfluous recovery actions, excessive resource utilisation, and operational disruptions. Policy optimisation is attained using Multi-Agent Deep Reinforcement Learning (MADRL) employing a centralised training and decentralised execution approach. Candidate learning algorithms encompass Multi-Agent Deep Q Networks (MADQN), Multi-Agent Proximal Policy Optimisation (MAPPO), and Multi-Agent Actor-Critic (MAAC), facilitating intelligent agents to collaboratively acquire adaptive defensive strategies while preserving autonomous decision-making at individual endpoints.

Layer 4 – Autonomous Response and Self-Recovery Layer

Subsequent to policy selection, the Autonomous Response and Self-Healing Layer converts intelligent decisions into actionable defensive measures via the Autonomous Response Agent (ARA). The ARA autonomously performs containment and remediation tasks, including terminating malicious processes, isolating compromised endpoints, blocking harmful network connections, suspending affected user accounts, quarantining malicious files, restoring interrupted services, and deploying relevant security patches. In contrast to traditional Endpoint Detection and Response systems that terminate the response process post-containment, the proposed framework has a specialised Recovery Agent (RA) tasked with reinstating endpoints to secure operating states. The Recovery Agent executes system rollback, configuration restoration, integrity verification, automated patch deployment, and ongoing post-recovery monitoring to guarantee that compromised endpoints return to secure operation while reducing the risk of reinfection. By including autonomous recovery into the cyber-defence lifecycle, the framework converts endpoint protection from a reactive detection-and-response model to a proactive, self-healing cyber resilience architecture.

Layer 5 – Knowledge Coordination Layer

The Knowledge Coordination Layer serves as the collaborative communication foundation that links all intelligent agents within the framework. Instead of perpetually transmitting raw telemetry, agents provide contextual information such as threat intelligence, acquired reinforcement learning policies, forensic evidence, attack context, historical reaction decisions, and recovery outcomes.

This collaborative knowledge repository enhances collective intelligence by allowing agents to leverage the experiences of other endpoints, while concurrently minimising communication costs and maintaining endpoint autonomy. Additionally, the layer coordinates policy updates produced using reinforcement learning, guaranteeing that insights gained from a single security incident are shared across the organization to enhance future detection, investigation, response, and recovery strategies.

Layer 6 – Continuous Learning Layer

The Continuous Learning Layer finalises the autonomous cyber-defence lifecycle by allowing the framework to enhance its performance through accumulated operational experience. Each resolved cybersecurity issue is regarded as a learning opportunity from which useful insights are derived. Subsequent to each incidence, the efficacy of the chosen response is assessed, reinforcement learning rewards are computed, defensive strategies are revised, and newly gained knowledge is synchronised across collaborating agents. This ongoing feedback system enables behavioural detection models, decision-making policies, and autonomous response methods to develop incrementally as new assault scenarios arise. Thus, the suggested approach demonstrates adaptive intelligence that perpetually improves its defensive capabilities, facilitating more precise threat detection, enhanced autonomous reaction, and increased organisational cyber resilience over time.

4.2 Autonomous Multi-Agent Interaction Model

The suggested Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR) framework conceptualises business endpoint protection as a cooperative ecosystem of intelligent software agents that function autonomously while concurrently enhancing enterprise-wide collective learning. Each secured endpoint contains a localised assembly of specialised intelligent agents that can autonomously monitor, analyse, and respond to security incidents while sharing contextual information with other agents via a decentralised coordination system. The interaction model adheres to a perpetual cyber-defence lifecycle commencing with the procurement of endpoint telemetry, which offers extensive insight into endpoint activity. The gathered telemetry is then subjected to behavioural anomaly detection to discover departures from standard operational patterns.

Upon detecting unusual behaviour, the framework conducts threat validation to differentiate authentic cyber threats from benign activities prior to commencing an automated forensic investigation. At this stage, contextual evidence is gathered and processed to facilitate informed decision-making. The resultant data is subsequently assessed by the reinforcement learning decision engine, which establishes the ideal defensive strategy based on the prevailing security scenario, prior learning experiences, organisational policies, and projected long-term rewards. The chosen policy is then implemented via autonomous reaction mechanisms, succeeded by self-healing endpoint recovery processes aimed at restoring compromised systems to trusted operational conditions. Following the successful resolution of the incident, the consolidated operational experience is coordinated among collaborating agents to promote ongoing organisational learning and enhance future defensive efficacy. Intelligent agents communicate via an asynchronous event-driven messaging architecture, facilitating low-latency information transmission, rapid incident response, and high scalability, while maintaining endpoint autonomy and reducing communication overhead.

5. CONCLUSION

This study proposes a conceptual Autonomous Self-Healing Endpoint Detection and Response (ASH-EDR) Framework that integrates Multi-Agent Artificial Intelligence, Multi-Agent Deep Reinforcement Learning, autonomous response, and recovery mechanisms into a unified endpoint security architecture. The framework addresses identified gaps in existing approaches by conceptualizing how distributed intelligence, adaptive decision-making, and self-healing capabilities could support the endpoint incident response lifecycle.

Contribution to Research

The study contributes to research by providing an architectural reference for investigating the integration of multi-agent coordination, adaptive defensive policies, and autonomous recovery.

Contribution to Practice

For industry practice, it offers design considerations for developing endpoint security systems with reduced dependence on manual intervention.

Contribution to Policy

For policy and governance, it highlights the need for explainability, accountability, privacy protection, human oversight, and safeguards governing autonomous security actions.

As the framework remains conceptual, its anticipated benefits including improved detection, response efficiency, and cyber resilience have not yet been empirically demonstrated. Its contribution therefore lies in architectural synthesis and the formulation of a foundation for further investigation, rather than evidence of operational effectiveness. Future research should implement a functional prototype and evaluate it experimentally against appropriate baseline EDR approaches. Evaluation should measure detection accuracy, false-positive rates, mean time to detect and respond, recovery success, resource consumption, scalability, and human intervention requirements. Reproducible empirical results should establish whether ASH-EDR delivers measurable performance improvements over existing approaches. Further investigation should also address adversarial robustness, explainable artificial intelligence, secure agent coordination, and the reliability of autonomous recovery. Ultimately, the significance of ASH-EDR will depend on whether subsequent implementation and empirical validation demonstrate that its proposed architectural integration offers practical and measurable advantages for endpoint security and organizational resilience.

REFERENCES

1. Abouhawwash, M. (2024). Innovations in cyber defense with deep reinforcement learning: A concise and contemporary review. *Artificial Intelligence in Cybersecurity*. https://doi.org/10.61356/j.aics.2024.1298
2. Carvalho Junior, H. (2024). HookChain: A new perspective for bypassing EDR solutions. *arXiv*. https://doi.org/10.48550/arXiv.2404.16856
3. Finistrella, S., Mariani, S., & Zambonelli, F. (2024). Multi-agent reinforcement learning for cybersecurity: Approaches and challenges. In *Proceedings of the 25th Workshop "From Objects to Agents" (WOA 2024)* (pp. 103–118).

4. Finistrella, S., Mariani, S., & Zambonelli, F. (2025). Multi-agent reinforcement learning for cybersecurity: Classification and survey. *Intelligent Systems with Applications, 26*, 200495. https://doi.org/10.1016/j.iswa.2025.200495
5. Foley, M., Hicks, C., Highnam, K., & Mavroudis, V. (2024). Autonomous network defence using reinforcement learning. *arXiv*. https://arxiv.org/abs/2409.18197
6. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105.
7. Hevner, A. R., Parsons, J., Brendel, A. B., Lukyanenko, R., Tiefenbeck, V., Tremblay, M. C., & vom Brocke, J. (2024). Transparency in design science research. *Decision Support Systems, 182*, 114236.
8. Kaur, H., Dharani, S. S. L., Paul, T., Thakur, R. K., Reddy, K. V. K., Mahato, J., & Kaviti, N. (2024). Evolution of Endpoint Detection and Response (EDR) in cyber security: A comprehensive review. *E3S Web of Conferences, 556*, 01006. https://doi.org/10.1051/e3sconf/202455601006
9. Landolt, C. R., Würsch, C., Meier, R., Mermoud, A., & Jang-Jaccard, J. (2025). Multi-Agent Reinforcement Learning in cybersecurity: From fundamentals to applications. *arXiv*.
10. Mavroudis, V., Palmer, G., Farmer, S., Whitehead, K. S., Foster, D., Price, A., Miles, I., Caron, A., & Pasteris, S. (2025). Guidelines for applying RL and MARL in cybersecurity applications. *arXiv*.
11. Mohamed, N. (2025). Artificial intelligence and machine learning in cybersecurity: A deep dive into state-of-the-art techniques and future paradigms. *Knowledge and Information Systems, 67*, 6969–7055. https://doi.org/10.1007/s10115-025-02429-y
12. Moore, A., Burke, A., Foley, M., Knack, A., Hicks, C., & Mavroudis, V. (2025). *A fundamental research plan for autonomous cyber defence*. Centre for Emerging Technology and Security.
13. Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77.
14. Pell Reddy, R. R. (2024). Enhancing endpoint security through collaborative zero-trust integration: A multi-agent approach. *International Journal of Computer Trends and Technology, 72*(8), 86–90. https://doi.org/10.14445/22312803/IJCTT-V72I8P112
15. Secure autonomous cyber defense with LLM agents: A systematic review of autonomy, tool-augmented reasoning, and governance constraints. (2026). *Computers & Electrical Engineering*.
16. Sharma, A., Patel, R., & Chen, Y. (2025). AI-augmented SOC: A survey of LLMs and agents for security automation. *Journal of Cybersecurity and Privacy, 5*(4), 95. https://doi.org/10.3390/jcp5040095
17. Shen, X., Li, Z., Burleigh, G., Wang, L., & Chen, Y. (2024). Decoding the MITRE Engenuity ATT&CK Enterprise Evaluation: An analysis of EDR performance in real-world environments. *arXiv*. https://doi.org/10.48550/arXiv.2401.15878
18. Singh, A. V., Rathbun, E., Graham, E., Oakley, L., Boboila, S., Chin, P., & Oprea, A. (2025). *Hierarchical Multi-agent Reinforcement Learning for cyber network defense*. *Proceedings of the Reinforcement Learning Conference (RLC 2025)*.
19. Tuunanen, T., Winter, R., & vom Brocke, J. (2024). Dealing with complexity in design science research: A methodology using design echelons. *MIS Quarterly, 48*(2), 427–458.
20. Verma, D., Terziyan, V., Tuunanen, T., & Shukla, A. K. (2026). Toward an artifact that designs itself: Generative design science research approach. *AI and Ethics, 6*, Article 104.

21. Vyas, S., Bolton, A., Mavroudis, V., & Burnap, P. (2025). Towards the deployment of realistic autonomous cyber network defence: A systematic review. *ACM Computing Surveys, 58*(1), Article 5. https://doi.org/10.1145/3729213
22. Zhang, J., Bu, H., Wen, H., Liu, Y., Fei, H., Xi, R., & Li, L. (2026). Agentic AI and large language models for autonomous IoT cybersecurity: A systematic survey, taxonomy, and research roadmap. *Electronics, 15*(12), 2740.
23. Zhu, C., Dastani, M., & Wang, S. (2024). A survey of multi-agent deep reinforcement learning with communication. *Autonomous Agents and Multi-Agent Systems, 38*, Article 4. https://doi.org/10.1007/s10458-023-09633-6