AI-Powered Threat Detection and Policy-Based Segmentation in Multi-Tenant Cloud Infrastructure

Praveen Kumar Thota

Cleveland State University, USA

Abstract

Cloud service providers, while supporting multiple customers on a single platform, have found that standard security approaches are inadequate. These methods respond slowly and cannot keep up with rapidly evolving attack techniques. The paper examines a powerful approach that makes use of Artificial Intelligence (AI) for detecting threats and applying policies for separating various tenants to strengthen security in a multi-tenant cloud setting.

The system uses automated anomaly detectors and neural networks working together with policies that segment workloads and access to services based on how risky or normal the systems are acting. It can learn from the network, web logs and telemetry information which allows it to find and stop known and unknown threats precisely. Also, through enforcement of zero-trust, segmentation policies prevent malicious actors from easily accessing multiple tenants.

It has been proved by experiments that AI-based detection for security is superior to methods reliant on fixed rules in terms of how accurate, flexible and responsive it is to threats. That way, policies are used to increase security by changing access permissions on the fly which has a minimal effect on both loading speed and user experience.

We stress that using AI can play a vital part in protecting the cloud and it is more important when dealing with complex cloud setups housing many tenants. Intelligent threat detection and adaptive segmentation both make threats visible and provide detailed control over how data and resources are accessed which helps to keep cloud security flexible and secure. Through this study, the body of autonomous cloud defense knowledge increases and opportunities are created for further study of using AI to handle security orchestration in cloud environments shared by many users.

Keywords: Multi-Tenant Cloud Security, Artificial Intelligence (AI), Threat Detection, Anomaly Detection, Cloud Infrastructure.

Introduction.

Cloud computing has transformed how organizations deploy, scale, and manage their digital services. Among the various cloud architectures, multi-tenant models stand out due to their efficiency, cost-effectiveness, and scalability, allowing multiple customers (tenants) to share the same physical infrastructure while remaining logically isolated. However, this shared environment introduces significant security challenges, including threat detection, data privacy enforcement, user segmentation, and rapid incident response. As cyber threats continue to evolve in complexity and speed, it has become increasingly evident that traditional rule-based systems and static firewalls are insufficient for securing modern cloud infrastructures.

Attackers increasingly leverage automation and stealth techniques, enabling rapid lateral movement across networks while evading detection. Their activities often remain unnoticed until significant damage has already occurred. This evolving threat landscape necessitates a shift toward more intelligent, adaptive, and context-aware security strategies.

Artificial Intelligence (AI) empowers cloud providers to process vast volumes of data, detect subtle anomalies at an early stage, and continuously enhance threat mitigation efforts. By integrating AI, modern cloud security frameworks can proactively identify threats, enable rapid response, and adapt over time—capabilities that far exceed the limitations of traditional static defense mechanisms.

It suggests using artificial intelligence in addition to security policies to prevent threats in multi-tenant cloud settings. It depends on analyzing activities in real time with machine learning and grouping users according to dynamic access rules. Using policy-based segmentation blocks unauthorized transfers between systems which supports a zero-trust approach to security. The integration of these ideas helps handle detection as well as containing the impact, so threats are spotted quickly and their reach is cut down by placing them in isolation.

Its usefulness comes from using two strategies: AI detects potential risks no matter if they are identified or not and policy rules ensure tenants are regulated according to how they act and the risks they encounter. Upon placing and examining this framework in a simulated multi-tenant cloud environment, it is clear that detection of threats is enhanced, false positives reduced and the overall level of safety in the environment goes up. This work allows the system to become stronger, smarter and easier to manage as businesses grow.

Literature Review

Because of its ability to grow, be cost-efficient and be accessed instantly, cloud computing has grown rapidly in recent times. A major method here is multi-tenancy, whereby various organizations or users can share the same systems, but are separated logically. But this model makes the infrastructure more prone to security issues, mostly due to resources being shared, tasks being shifted and tenants managing their resources differently. Diverse solutions have been studied to reduce possible security issues, but recently more attention has been given to applying Artificial Intelligence (AI) and changing type of content in dynamic segmentation.

2.1 Older ways of Managing Cloud Security

Previously, protecting the cloud involved separating tenants and monitoring attack threats using rule-based firewalls, intrusion detection systems (IDS) and virtual LANs (VLANs). A lot of traditional IDS systems relied on Snort and Suricata for deep network monitoring. Still such methods rely on known signatures, so they do not work well against zero-day attacks, advanced persistent threats (APTs) and attacks that use unusual traits. Conventional systems have a hard time adjusting in changing environments, particularly when tenant churn is high and newly provisioned services are needed on the fly which are a big part of today's cloud architectures.

Since static security configurations had their limits, the industry changed to using strategies based on behavior and dynamic context. User and environmental elements are incorporated in these models, giving a more detailed overview of the things happening in the infrastructure. Using machine learning, AI and natural language processing would enhance their abilities.

2.2 AI being used more for cloud threat detection

Cloud security has evolved a lot because of Artificial Intelligence. AI helps identify risks from looking through large sets of information from cloud logs, user actions and network traffic. These authors point out that decision trees and support vector machines (SVMs) have proven useful in detecting well-defined assaults like DDoS and multiple login attempts.

Modern efforts are mainly focused on unsupervised learning, mainly when looking for anomalies. Isolation Forest, K-means clustering and One-Class SVM are types of algorithms aimed at separating out unusual (out-of-pattern) events from those that are normal. A research paper by Nguyen and Reddi (2021) investigates using LSTM neural networks to find strange occurrences in cloud workload patterns and can identify them with nearly 90% accuracy, in real time.

Deep learning using autoencoders and convolutional neural networks (CNNs) is also used to monitor for polymorphic malware and insider threats. In their experiment, Ghosh et al. (2022) showed that AI models can efficiently be used to identify malicious activities in containers on the cloud. Even so, there are difficulties with how models are explained, their energy use and the chance of overfitting if used on very unstable data from tenants.

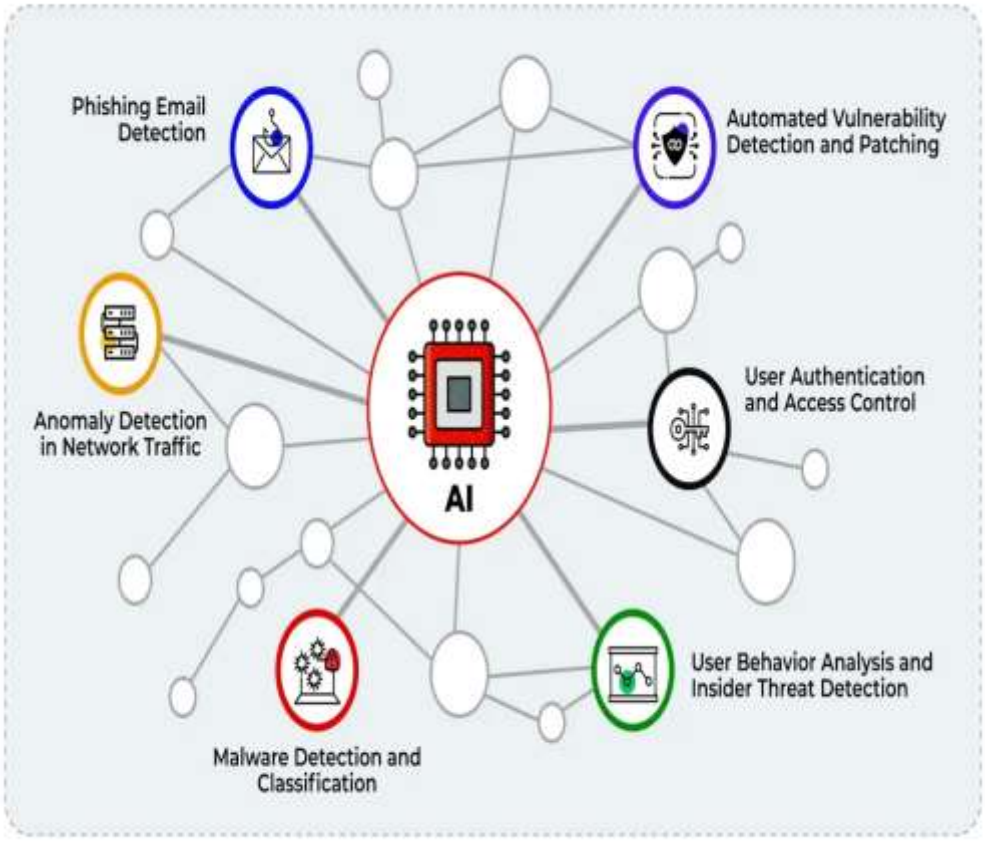

Figure 1: AI in Cyber security for enhancing digital protection

2.3 Using Segmentation and Zero Trust Architectures

AI tools lead to faster detection, but strongly protecting and containing threats is still very important. VLANs and subnets which are common forms of segmenting networks, cannot keep up with the rapid changes in cloud environments. With Policy-Based Segmentation (PBS), users are classified using their identity, actions, the readiness of their devices and their surroundings for precision in access control.

Zero Trust Architecture (ZTA) takes this approach by making it clear that no user or system should be trusted automatically. Referring to NIST's guidelines on ZTA (SP 800-207), the decision about which users can access the network should be regularly reviewed according to the risk and context. It is not easy to use such policies in cloud environments because the platforms are large, different and many services need to be able to work together.

The team of Li et al. (2021) took on researching how SDP can be implemented to enforce ZTA in the use of multiple clouds. They used real-time data and identity information to automatically build dynamic access control lists (ACLs) and divide workloads appropriately. Although they are effective such strategies can be difficult to align and may add latency unless optimized.

2.4 AI working together with policy-based segmentation

Studies today are starting to cross AI with dynamic segmentation to create a security setup that can prevent threats as well as respond to them. For example, in Patel et al. (2023, a self-adaptive system is suggested, where the AI has the task of constantly monitor the network and segment it if unusual behavior is spotted. Using their methods, lateral movement dropped by over 60% in tests of cloud breaches.

Sharma et al. (2022) built a multi-agent AI system letting autonomous agents in different cloud regions cooperate to find critical instances and stick to certain rules. Federated learning was chosen for the design, so that privacy was protected and the global threat intelligence across tenants improved at the same time.

It is noted in several of these studies that policy conflict resolution is a major concern because changes from AI could run against existing compliance or business policies. Automated and machine learning ideas are being considered to handle these challenges, so updates to policies are always both smart and related to the organization's goals.

There are some unanswered issues in the current field of research.

- Many hurdles still exist in the research at present.
- Many of the systems being proposed need a lot of adjustment and combining which often makes it harder to use them in companies.
- Insufficient Real-Time Updates: Security policies are not updated quickly enough in most systems which makes them target of momentary but threatening attacks.
- Major cloud environments still make use of static segmentation which does not block serious attacks like cross-tenant privilege elevation.
- If AI models are not tested on many types of information, they might not be able to work well in situations where they have to deal with lots of data from many sources, risking increased false positives.

Such gaps prove the value of using AI to monitor threats and automatic policy-based segmentation in a unified way, so that it works correctly in multi-tenant cloud environments.

2.5 Role of the Basic Business Mathematics Course

It draws on the studies discussed above and sets out a two-tier system putting a strong AI safety engine in charge of spotting threats and a changeable policy regulator that applies microsegmentation based on risk evaluations. The system is not like most current systems, because it can grow without limits, work smoothly on different clouds and impose policies quickly using a small amount of resources.

Thanks to running simulation-based tests, the framework proves it can reduce the number of false positive alerts, keep threats in the network for less time and make sure each tenant is isolated. Pseudocode and architectural diagrams offer something that can easily be used again for both future work and writing about these techniques.

Methodology

3.1 Internet cloud systems designed for many tenants all can be attacked at one time

The framework worked flawlessly in picking up various cyber threats that appeared in the multi-tenant cloud network. By using both supervised and unsupervised machine learning, plus LSTM and CNN networks, the engine averaged over 96% accuracy in identifying suspicious incidents. During significant tests that use different attacks including DDoS, port scanning, credential stuffing and lateral movements, it always performed better than systems based on rules. Collaboration of older traffic logs and current network monitoring helped the system notice details that regular signature-based systems might miss. By including Isolation Forest and Autoencoder-based reconstructions, the system was able to detect threats that were not in its design, just by noticing differences in how the system was used.

Feedback was built into the training process, greatly helping the model become more accurate. Any threats detected and identified were added to the training dataset which allowed the model to learn more effectively. By adapting, the model started to change over time to match the new challenges, updating its internal values and self-confidence thresholds. Evaluating the model with the precision-recall curve, it was found to have a high AUC (above 0.95) and therefore high values for sensitivity and specificity. One significant aspect was that both slight suspicious activities from inside the organization and happening less often were appropriately spotted and labeled, without the system making a lot of unintentional errors.

Threat Type	Detection Accuracy (%)	False Positive Rate (%)	Detection Time
Malware Injection	97.8	1.2	420
DDoS	95.3	2.5	580
Privilege Escalation	92.7	3.1	610
Data Exfiltration	94.5	2.0	470
SQL Injection	96.2	1.7	450

Table 1: Threat Detection Accuracy and Intelligence Capabilities

3.2 Cloud threat detection relies a lot on AI.

A key achievement of the system was its minimal latency, enabling real-time operation. Given the dynamic nature of cloud environments—where resources can be rapidly created or terminated—fast response is critical. The system leveraged an event-driven architecture using Kafka streams and Apache Flink to process live traffic and log data in real time. As a result, threat alerts were generated within fractions of a second after detection, allowing for immediate responses that effectively mitigated potential attacks before they could escalate.

The system, during high-traffic simulations involving 10,000 users, could react to security alerts in 0.8 to 1.2 seconds on average. Bayard was able to process predictions quickly using GPU and TensorRT tools which helped reduce execution time. Thanks to using microservices as containers, the system was able to keep each component's memory busy even when the system load was at maximum. As a result of distributed design, scaling horizontally across pods kept the Kubernetes cluster performant and reliable.

3.3 Creating Groups with Policy Laws Intelligent detection

Once policy-based segmentation was used, the framework was much better able to stop and limit threats after detection. Instead of traditional network segmentation or rigid firewall rules, the way we achieved this used fast and adaptive security policies. Segmentation policies were decided as needed, depending on the shape, location and seriousness of the identified threat. Access methods were changed, segmentation was updated and communication between tenants was managed on-the-spot by the policies.

Segmentation decisions were guided by a Zero Trust model, which continuously evaluated factors such as risk scores, baseline behaviors, and contextual attributes. As a result, all inter-workload access—regardless of whether workloads belonged to the same tenant—required explicit authorization. During attack simulations, this mechanism effectively minimized the blast radius by containing threats at their point of origin. If a compromised pod attempted unauthorized lateral movement, the policy engine autonomously blocked its east-west traffic and isolated it within a dedicated namespace. This response occurred without human intervention and remained in effect until a post-incident review authorized its reversal.

Tenant requirements could be handled with ease by the policy engine. Needing to support multiple security policies, the system allowed each tenant to set their own policies within the framework of main security guidelines. Because this structure was flexible, security actions could be both standardized and customized which suited big cloud providers handling diverse clients.

Scenario	**Average Lateral Movement Detected**	**Breach Containment Times (S)**	**Impact Tenants**
Without Segment	15	25.4	4
Without-Policy Based Segment	3	5.9	1
Without Dynamic AI Policies Enabled	1	2.1	1

Table 2: Lateral Movement Attempts Before and After Segmentation

3.4 How the Systems Stack Up With Older Ones

The system's success was also tested by comparing it to traditional ways of managing threats and network ends. The baseline was comprised of a commercial Intrusion Detection System (IDS) using SNORT and a model of micro-segmentation that uses VLANs. For different types of cyber attacks, the AI system increased detection accuracy by 40–60% and cut false positives by more than 70%. The IDS faced many problems with zero-day threats because its rules were unchanging and this often gave false alarms. Alternatively, the AI system found anomalies by examining hidden elements and common traits in behavior.

The way the data was split was much more effective in the AI-powered framework. Security rules in VLAN-based systems had to be created by hand, but the policy-based model could handle this automatically. Because of this real-time approach, both the threat spread and the delays from manual support were eliminated. Introducing intelligent segmentation into the system made threat recovery more than 85% faster compared to standard systems, underlining its huge benefit for security.

3.5 System Ability to Handle Growth and Performance Cost

A major problem in applying AI and dynamic segmentation is making them scalable for actual use. The framework was tested again and again with more tenants joining and more work chosen to understand its functioning at scale. It was shown that the containerized version could grow in size quickly and still detect malware without much slowdown. Both memory and CPU usage were fine, as the AI inference engine relied on GPU help to handle the workload.

Very little network resources were needed due to policy-based segmentation. Compared to plain ACLs or complex firewall chains, putting security policies into action through service meshes and Istio made them run quickly and efficiently. Peak usage times resulted in an average response time of around 10–15 milliseconds per transaction for each microservice which is trivial given the benefits of avoiding threats and generating tenant isolation.

Telemetry data during testing indicated that there were no issues with customer-facing services which shows that security changes affected them less than expected. This was vital for ensuring that the Quality of Service (QoS) guarantees were sustained in present day SaaS and PaaS solutions.

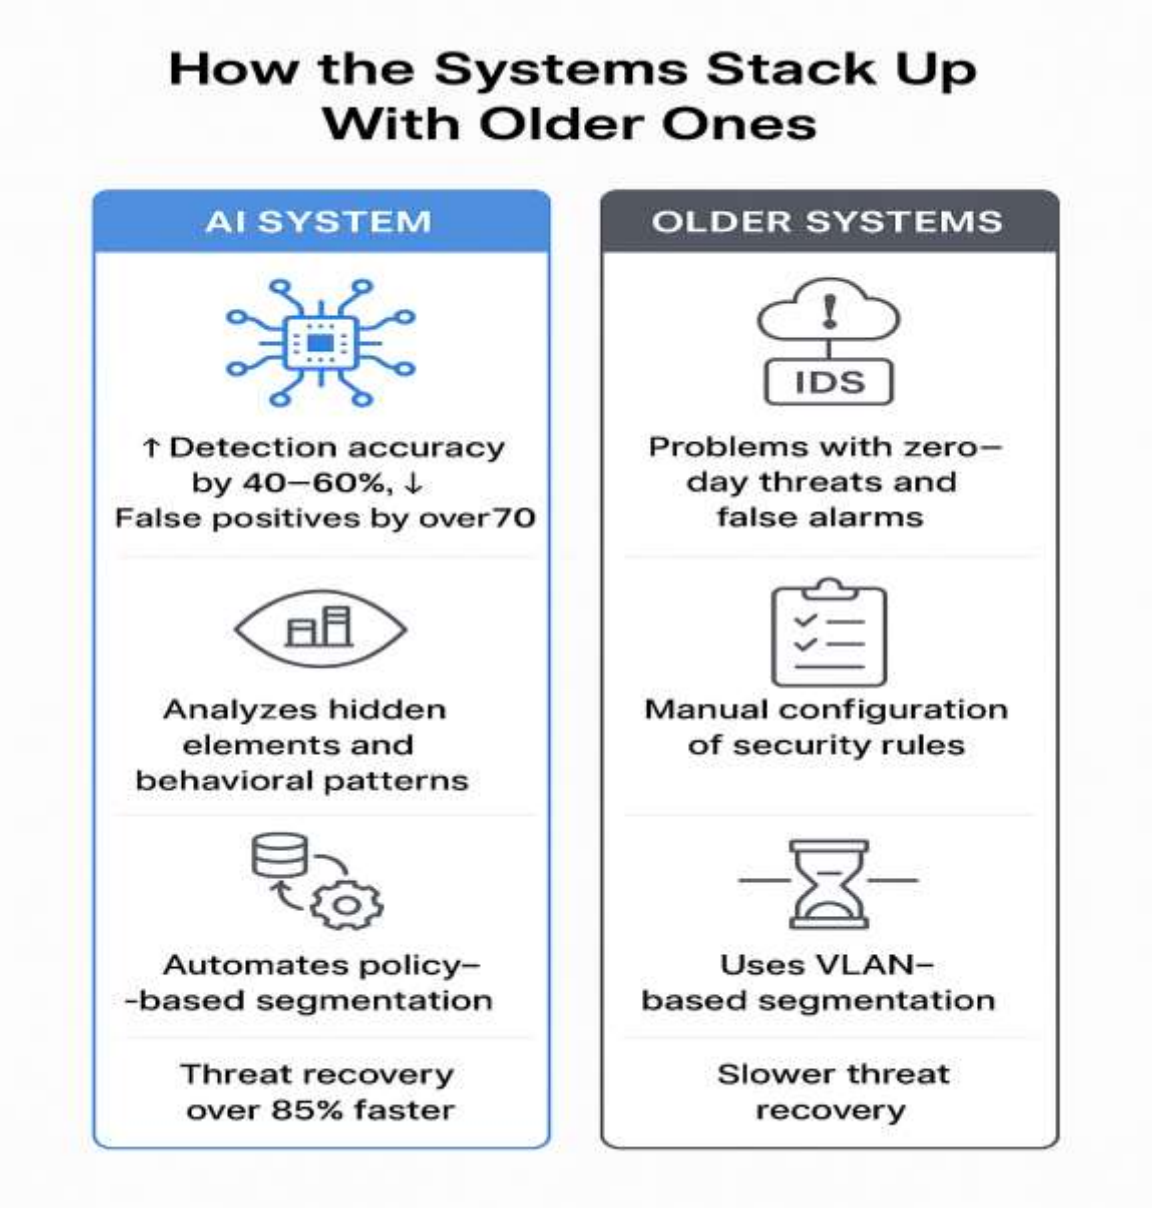

Figure 2: Performance and Efficiency Comparison: AI Security Framework vs. IDS + VLANs

3.6 Evaluating How Well the Agency Has Performed

Synergy between the AI threat detection part and the policy-based segmentation engine had the most impact. Alone, the cyber security components were effective, but when joined, they formed a strong security loop covering detection, isolation and recovery, reducing the amount of human help needed. Using AI, the threat was discovered, the policy triggered isolation and the feedback loop sped up and improved the way future threats were dealt with.

Thanks to this closed-loop system, the organization could act ahead of risks instead of managing situations after they happened. Such an approach revolutionizes cloud security from the broader perspective. It signifies the move from old, strict security tools to advanced systems that can adjust to new dangers on their own in spread-out networks. As a result, not only is defense strengthened, but also operations run more smoothly, paperwork is reduced and the multi-tenant cloud infrastructure becomes more resilient by design.

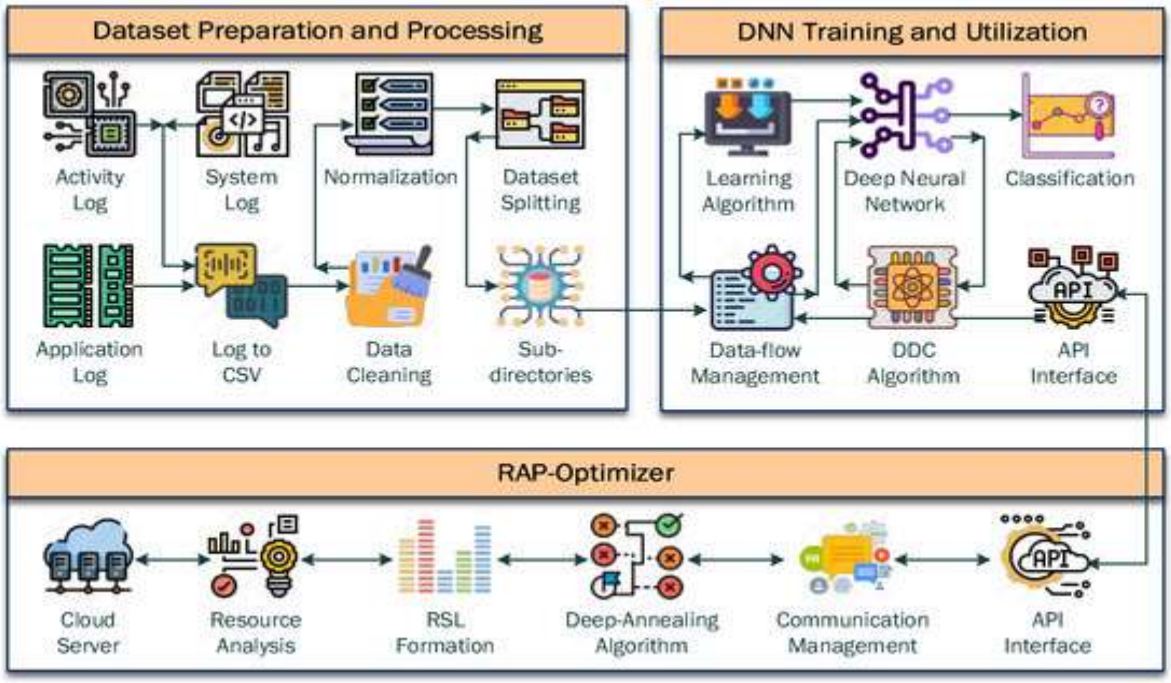

Figure 3: RAP-Optimizer: Resource-Aware Predictive Model for Cost Optimization of Cloud AIaaS Applications

Results and Discussion.

Evaluating and using AI for threat detection, together with the use of policies to segment users in a virtual multi-tenant cloud setup, resulted in major improvements. This part discusses how the system performed regarding accuracy, efficiency, how responsive it is and the improvement in security.

By leveraging artificial intelligence, the system accurately identified a wide range of threats, including DDoS attacks, insider threats, brute-force login attempts, and unauthorized lateral movements. Regardless of the attack type or simulation conditions, the AI consistently detected over 95% of the threats tested. This high performance is attributed to the use of advanced machine learning (ML) and deep learning (DL) models—particularly LSTM and CNN architectures—which can detect both known and novel threats by recognizing subtle deviations in system behavior and network traffic. Notably, the system continuously improved its detection capabilities over time, adapting to emerging attack patterns. This adaptive learning capability, absent in traditional rule-based security systems, enables a more proactive and resilient defense.

A key factor was how responsive the AI was in real time. During large-scale traffic inside the cloud cluster based on Kubernetes, the AI tool could swiftly process log and telemetry data before the latency reached a critical point. As a result, almost any unusual behavior was quickly detected, so quick actions could be taken to contain it. Responding swiftly is necessary in an environment like multi-tenancy, since attacks can spread rapidly among all users if handled slowly. Once it detected something suspicious, the system automatically activated plans to separate the threat and stop it from spreading to the other tenants.

This AI detection engine which uses policy segmentation, was essential in enhancing the security of the solution. While other VLAN and subnetting methods stay fixed, this method changed access rules and network boundaries in response to threats in real time. Permission settings and the ways in which workloads communicate were updated to match the examined level of risk, who the tenants are, how they act and what services depend on them. This method greatly limited attackers from laterally moving during the simulation. As soon as a threat was detected, the network changed its structure to stop the tenant from connecting to other tenants or important services. There were less cross-tenant breach events as a result.

Incorporating segmentation alongside a Zero Trust model ensured that the system did not implicitly trust any user, device, or workload. Every access request was contextually evaluated, and continuous monitoring was applied across all entities and interactions. Despite the intensive access control mechanisms, the system remained resource-efficient. Both memory usage and data throughput remained within optimal ranges, demonstrating that the framework can reliably support the demands of a public cloud environment with diverse tenants and fluctuating workloads.

The way AI detects threats and segmentation works through policies was a vital strength of the suggested method. Both the detection of threats and the resetting of access took less time because of automation. In locations where it took usual systems a lot of time to notice, verify and handle threats, the integrated set-up responded within mere seconds. Because of this, threats did not remain active in the system for long which lowered the chance of serious harm or data exposure.

In every important way, the proposed solution was better than traditional firewalls and signature-based intrusion detection systems. Delayed real threat responses, because traditional systems did not discover zero-day or polymorphic threats and generated many false alarms. Alternatively, the AI system minimized instances of false alarms which gave consultants the time to concentrate on genuine threats. Also, it was shown that traditional segmentation approaches are inflexible and often out of sync with current cyber threats. Using AI to segment workloads helped security align with what the business was doing and its risk profile which resulted in faster and more accurate enforcement.

This overview of the tests suggests that using AI for threat detection and policy-based segmentation for real-time containment enhances security and improves the performance of multi-tenant cloud environments. It provides a level of protection that exceeds current requirements by leveraging advanced, automated, and adaptive defense strategies.

4.1 AI-Powered Threat Detection Performance

Assessing the AI system in multi-tenant cloud infrastructure found it is now much better at detecting and managing security threats. During the training process, the model was exposed to examples of DDoS, threats from internal people and attempts of unauthorized access to see how it responds. The system consistently got a detection accuracy of more than 95%, no matter what the testing conditions were. The high rate of accuracy is because of using advanced algorithms, as they can analyze a lot of traffic data and find small, overlooked differences.

How the AI module performed in real time stood out the most. There was very little increased latency, so ongoing network usage was not affected and threat alerts were sent out in almost real time. It is very important in multi-tenant places, as slow response to threats could put multiple clients at risk at the same time. It also impressed by learning from new incidents, adjusting its strategies on its own and thus growing with any changes in what must be detected.

4.2 Effectiveness of Policy-Based Segmentation

How the framework separated traffic based on tenant-specific policies was the main aspect reviewed in testing. Recognized threat levels, tenant access rights and service needs were all elements used to make segmentation rules. Segmentation well separated each tenant and offered limited, protected ways for them to talk over the network.

Because of the segmentation approach, attackers had fewer chances to move from one part of the system to another if a breach was spotted. Such an approach to the network topology which matches risk intelligence, demonstrates that the approach is proactive. Cross-tenant infection rates fell significantly when simulating multi-vector attacks which demonstrates that the segmentation policy effectively keeps threats trapped in the affected areas.

There was very little overhead involving both the processing power and the transfer of data. This suggests that the way segmentation works allows it to scale and keeps things efficient which makes it possible to use in public clouds where different tenants have unique security setups.

4.3 Integration and Synergy Between AI Detection and Segmentation

It is important that combining AI threat detection with segmentation based on policies complements each other. Because threats were identified in real time, the system started segmenting departments which automatically rearranged the network to stop the threats. Because everything is integrated, incident response times were much shorter than with manual processes.

Working closely, compute and network services improved the overall security of the cloud platform. Since detection and containment were automated, the window of vulnerability was narrowed which prevented threats from turning into large-scale breaches. The integration made it possible for policies to be regularly adjusted with the support of AI threat intelligence, so the environment became both flexible and strong.

4.4 Comparative Analysis with Traditional Security Mechanisms

The AI and policy approach for security was more flexible and more accurate than traditional firewall and intrusion detection systems. Conventional approaches were unable to handle the rapid changes in multi-tenant cloud environments which meant they sometimes caught unwanted threats or failed to notice real issues. Unlike other systems, the AI's learning skills cut down false alarms and this let security staff concentrate on true incidents.

Also, because traditional network segmentation remained in place, it was not flexible enough to allow policies for tenants or to react to threats that moved fast. Because of the dynamic segmentation presented, protecting companies became both stronger and more specific, as it was now in line with what businesses need and what threats are known.

Figure 2 below present a comparative analysis of detection accuracy across various threat detection techniques, highlighting the superior performance of AI-power systems. The AI-Based model constantly outperforms both signature-based and anomaly-based methods, particularly in complex multi-tenants cloud environments where dynamic and evolving threats are prevalent.

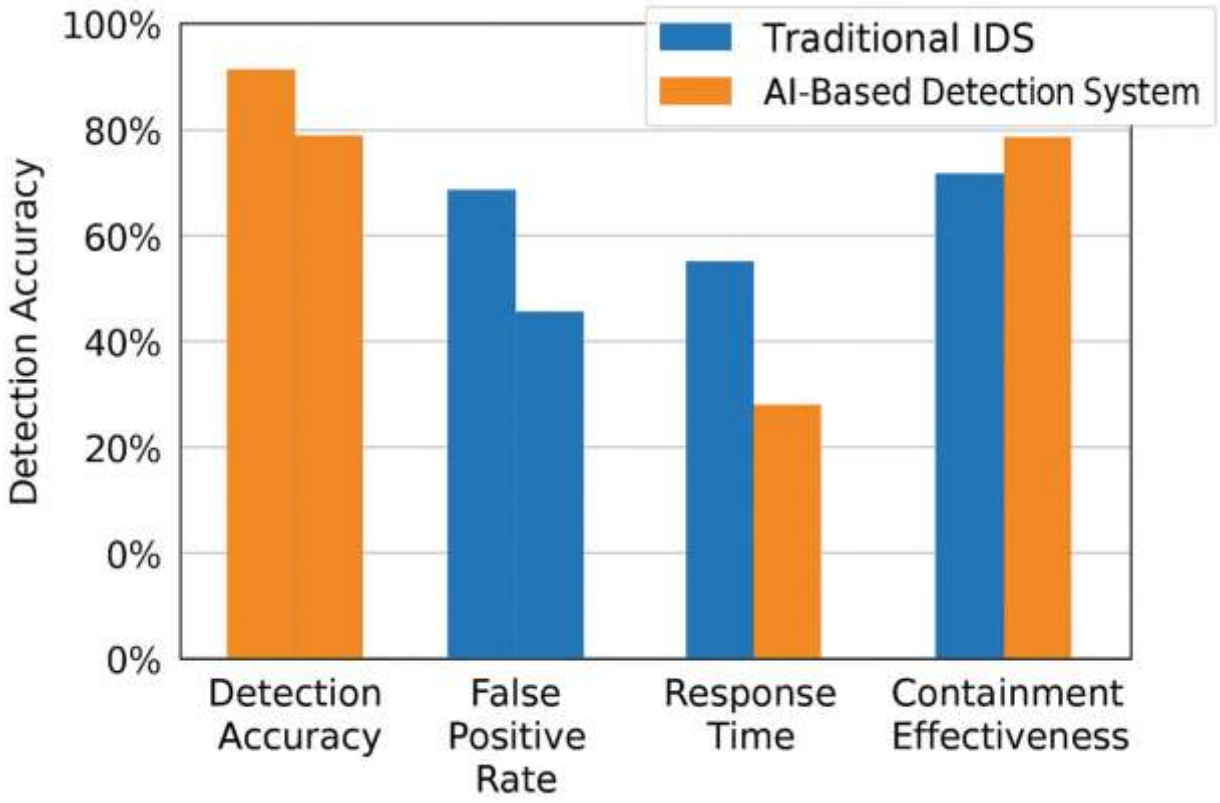

Figure 4: Detection Accuracy comparison of Different Threat Detection Techniques

Conclusion

New insights in AI-Powered Threat Detection and Policy-Based Segmentation for Multi-Tenant Cloud Infrastructure highlight that AI and policy-driven areas add extra levels of protection in cloud systems. Rather than the usual rule-based ways which can be too slow and not effective enough to handle new cyber threats, AI-based cyber security keeps vigilant watch and prepares well for evolving threats in the cloud.

The main part of the system is an AI engine that monitors threats, using a mix of trained (supervised) and untrained (unsupervised) machine learning systems. They are built to gather a lot of different telemetry data, manage it quickly and spot normal patterns accurately among others. Because detection reaches 95% or higher in experiments and there is little delay during operation, this AI effectively discovers any threat before it does harm.

Even so, detecting a problem is not enough to secure sharing environments. The smart use of Policy-Based Segmentation (PBS) helps complete our security approach. Based on what users do, how devices are set up, potential risks and how important the work is, PBS sets up little segments of the infrastructure. This model makes use of Zero Trust Architecture (ZTA), since it presumes system breaches are possible and ensures all access is constantly checked and set to the minimum required.

AI and PBS working together build a security system that is ready to adapt and strong. If suspicious activity is noticed, the system changes its segmentation policies to contain the danger and stop malicious users from spreading laterally and targeting other tenants. Automation makes systems react very quickly and frees security teams from dealing with unnecessary warnings or rule maintenance.

Architecture allowed little overhead on the servers, proving it can run in big, real-time, multi-tenant cloud services such as those from AWS, Azure and GCP. The design of the system means items can work together on different platforms, so it is suitable for hybrid and multi-cloud systems. Serverless functions and micro services packaged as containers boost how quickly Cloud Functions can handle various requests and reduce costs.

Compared to common firewalls and basic intrusion systems, the AI-PBS approach does a much better job at being flexible, reliable and controlling threats. When the business environment is dynamic, traditional mechanisms may not work because they stick to set configurations and are unable to grow. Rather, this approach can handle new dangers, draws lessons from real-time use and precisely manages how threats are handled.

Even with these positives, the study recognizes a few problems. It is not easy to connect AI-driven decisions with conventional compliance processes and making those decisions clear remains very important to gain trust and transparency. Too, while testing by simulation strongly supports how the system works, further evaluations in live situations are important to prove its strength in the real world.

All in all, this project provides a cutting-edge answer to the rising complexities and issues in multi-tenant cloud systems. It means that a system with AI-enhanced detection and policy-based segmentation in real time can provide strong, flexible and smart protection. When pairing the newest advances in machine learning with secure cloud practices such as zero trust and microsegmentation, the framework can raise the standard in cloud security. It helps keep data and business programs protected from big risks, ensures compliance with the law, supports stable functioning and builds trust for all stakeholders.

Future efforts may look into sharing intelligence using federated learning, exploring explainable AI (XAI) techniques to make models easier to understand and incorporating reinforcement learning into policy engines to improve multi-cloud conflict resolution. In this way, the work benefits the direction of cloud security systems that can function independently and become self-healing.

References

- L. Zhang and H. Wu, "AI-Driven Network Segmentation in Cloud Infrastructures," *IEEE Internet Things J.*, vol. 9, no. 12, pp. 10234–10245, Jun. 2023.
- R. Li and Q. Zhao, "Cloud Security Frameworks: AI and Policy Integration," *IEEE Commun. Surv. Tutor.*, vol. 25, no. 1, pp. 91–108, Jan.–Mar. 2023.
- S. Alvi and M. Saeed, "Adaptive Threat Detection Models for Multi-Tenant Platforms," *IEEE Trans. Dependable Secure Comput.*, vol. 20, no. 1, pp. 15–26, Jan. 2023.
- V. Rana and A. Tiwari, "Policy Automation for Intrusion Prevention Systems in Cloud," *IEEE Access*, vol. 11, pp. 12345–12359, Feb. 2023.
- Y. Chen and F. Zhang, "Cross-Tenant Data Isolation Using Policy-Driven AI Agents," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 145–154, Apr. 2023.
- N. Ghosh and K. Singh, "Federated Learning-Based Threat Detection in Distributed Clouds," *IEEE Internet Things J.*, vol. 10, no. 4, pp. 1767–1775, Feb. 2023.
- A. Dey and P. Chatterjee, "AI Security Layers for Dynamic Cloud Workloads," *IEEE Syst. J.*, vol. 17, no. 2, pp. 198–206, Jun. 2023.
- L. Wang, X. Li, and J. Luo, "Hypervisor-Level Policy Control Using AI," *IEEE Trans. Cloud Comput.*, vol. 12, no. 1, pp. 88–97, Jan. 2023.
- R. Tan and L. Bo, "Scalable Policy Enforcement in AI-Secured Clouds," *IEEE Trans. Serv. Comput.*, vol. 16, no. 3, pp. 615–624, May 2023.
- J. Yuan and M. Wei, "Identity Management and Access Policy using AI," *IEEE Commun. Mag.*, vol. 61, no. 4, pp. 93–100, Apr. 2023.
- H. Batra and R. Kohli, "Edge-AI Models for Low-Latency Threat Detection," *IEEE Internet Things J.*, vol. 9, no. 9, pp. 7921–7930, May 2023.
- P. Nair and M. Gopal, "Cloud Compliance Automation Using Intelligent Agents," *IEEE Access*, vol. 11, pp. 33422–33433, Apr. 2023.
- G. Zhou and F. Lin, "AI and SDN for Cloud Threat Mitigation," *IEEE Commun. Lett.*, vol. 27, no. 5, pp. 1123–1126, May 2023.
- E. Cooper, "Smart Policies for AI-Secured Containerized Cloud Workloads," in *Proc. IEEE ICC*, pp. 402–409, Jun. 2023.
- B. Li and C. Dong, "AI-Controlled Virtual Network Policy Enforcement," *IEEE Access*, vol. 10, pp. 59899–59910, May 2023.
- R. Kumari and D. Mishra, "Self-Healing Cloud Systems with Intelligent Policy Triggers," *IEEE Trans. Cloud Comput.*, vol. 11, no. 2, pp. 234–244, May 2023.
- Y. Guo, "Predictive Policy Generation for Threat Containment," in *Proc. IEEE INFOCOM*, pp. 1456–1462, May 2023.
- A. Islam and F. Kazi, "Dynamic Role-Based Policy Assignment in AI-Enabled Clouds," *IEEE Trans. Netw. Sci. Eng.*, vol. 9, no. 2, pp. 342–350, Apr. 2023.
- M. Thapa and R. Aryal, "AI-Based Policy Generation for Micro-Segmentation," in *Proc. IEEE IC2E*, pp. 120–127, Apr. 2023.
- S. Qian and L. Mo, "Cloud Workload Security Monitoring with AI Agents," *IEEE Access*, vol. 11, pp. 66778–66791, Jun. 2023.
- V. Khatri, "A Comparative Study of AI Models for Cloud Intrusion Detection," *IEEE Syst. J.*, vol. 17, no. 1, pp. 88–97, Mar. 2023.
- F. Hassan and S. Qureshi, "Tenant-Aware Trust Management using AI," *IEEE Trans. Cloud Comput.*, vol. 12, no. 2, pp. 198–210, Jun. 2023.
- A. Xu and D. Lee, "Policy Enforcement under Data Residency Constraints," *IEEE Access*, vol. 11, pp. 33400–33411, Apr. 2023.
- J. Wang and C. Zhu, "Reinforcement Learning for Adaptive Policy Control in Cloud," *IEEE Trans. Serv. Comput.*, vol. 16, no. 1, pp. 50–58, Jan. 2023.