AI-Based Adaptive Loyalty Program Vulnerability Detection in Hotel Reservation Systems

Deepak Kaul
Marriott International

Abstract: Loyalty programs have become a cornerstone of customer retention strategies in the hospitality industry, particularly for hotel chains. However, their widespread use makes them prime targets for cyberattacks, exposing personal data and system vulnerabilities. Traditional security solutions are often inadequate in defending against sophisticated, multi-platform cyber threats. This paper proposes an AI-based adaptive model using federated learning to detect and respond to vulnerabilities in hotel loyalty programs. Federated learning enables the aggregation of insights from multiple hotel chains without compromising customer privacy, allowing for real-time anomaly detection across distributed systems. The model identifies potential attacks by analyzing anomalous patterns and offering preemptive responses. This approach introduces a novel layer of cybersecurity for hotel loyalty programs by addressing cross-platform vulnerabilities. Our findings demonstrate that AI, when combined with federated learning, can significantly enhance the security of loyalty programs while preserving data privacy and system integrity.

Keywords: AI-based vulnerability detection, federated learning, loyalty programs, hotel reservation systems, cybersecurity, cross-platform attacks, anomaly detection.

1. INTRODUCTION

The proliferation of loyalty programs in the hospitality industry has brought about a new wave of security challenges. Hotels increasingly depend on these programs to maintain customer engagement and build long-term relationships, but they also create vulnerabilities that cybercriminals exploit. According to a 2022 report by Accenture, the global hotel loyalty program market is projected to grow by 7.6% annually, increasing the size of this lucrative but vulnerable sector. Cyberattacks on such programs can lead to significant financial losses and severe reputational damage for hotel chains. The Marriott International breach in 2018, where the personal information of up to 500 million guests was compromised, serves as a stark reminder of the risks inherent in such systems.

Traditional cybersecurity measures, while effective to some degree, struggle to keep up with the sophistication of modern attacks, especially those targeting multiple hotel chains or platforms simultaneously. With the rise of machine learning (ML) and artificial intelligence (AI) technologies, there is an opportunity to revolutionize the detection and prevention of such vulnerabilities. However, the challenge remains in protecting sensitive customer data while effectively analyzing system-wide behaviors.

This paper proposes an AI-based adaptive model for vulnerability detection in hotel loyalty programs, utilizing federated learning (FL). Federated learning allows the analysis of distributed data across multiple hotel chains without the need to centralize sensitive information, thereby maintaining user privacy while improving system security. By using AI to identify anomalous patterns that indicate potential cyberattacks, we introduce a method that addresses the growing concern of cross-platform vulnerabilities in loyalty programs.

The novelty of this approach lies in the application of federated learning to loyalty program cybersecurity-a domain that has rarely been explored in existing research. This method not only offers improved security but also reduces the risk of sensitive data leakage. Our study investigates the efficacy of this model through real-world data and simulations across multiple hotel chains. We aim to fill a significant gap in the literature by demonstrating the potential of AI and FL in protecting against multi-chain, cross-platform attacks in hotel loyalty programs.

2. LITERATURE REVIEW

Loyalty programs, by their very nature, create a large repository of customer data, which makes them attractive targets for cyberattacks. Previous studies have explored the vulnerabilities in loyalty programs, highlighting the growing concern over privacy breaches. For instance, Ghandour and Haider (2019)

examined loyalty program breaches across multiple industries and identified a pattern of escalating sophistication in cyberattacks targeting loyalty rewards. According to their findings, 61% of loyalty program breaches go unnoticed for extended periods, allowing attackers to siphon points, personal data, or both.

Federated learning, a novel paradigm in AI, has shown great promise in enhancing privacy-preserving data analysis. McMahan et al. (2017) introduced the concept of federated learning, which allows multiple systems to train models collaboratively without sharing the underlying data. This approach has proven effective in sectors such as healthcare and finance, but its application in cybersecurity, particularly for loyalty programs in hospitality, remains underexplored.

Several studies have examined AI's role in cybersecurity. Sahay et al. (2020) proposed an AI-driven framework for real-time detection of cyber threats using anomaly detection techniques. Their study indicated that machine learning models, when trained on large datasets, could predict cyberattacks with 87% accuracy. However, the challenge of securing loyalty programs across multiple platforms has been left largely unaddressed.

3. METHODOLOGY

3.1 AI-Based Adaptive Model for Vulnerability Detection

The proposed AI-based adaptive model uses anomaly detection algorithms to identify unusual patterns in loyalty program activities. The model is designed to be scalable, capable of being implemented across multiple hotel chains, and adaptive, meaning that it learns from each detected anomaly to improve its response over time.

Table 1 provides an overview of the system components involved in the adaptive model.

Component	Description
Data Collection	Transaction data from distributed hotel loyalty programs
Anomaly Detection	AI algorithms for identifying irregular patterns
Federated Learning	Distributed learning to protect sensitive customer information
Response Mechanism	Real-time preemptive responses to detected threats

3.2 Federated Learning for Cross-Platform Security

Federated learning allows hotel chains to train the AI model collaboratively without sharing raw data, ensuring that sensitive customer information remains on local servers. This approach is particularly useful in detecting attacks that span multiple hotel chains. Each hotel chain contributes to the model by analyzing local data and sending updates to a central server, where the aggregated model is continuously improved. The use of federated learning ensures compliance with data privacy laws such as GDPR, which is a significant concern for hotels operating in multiple jurisdictions.

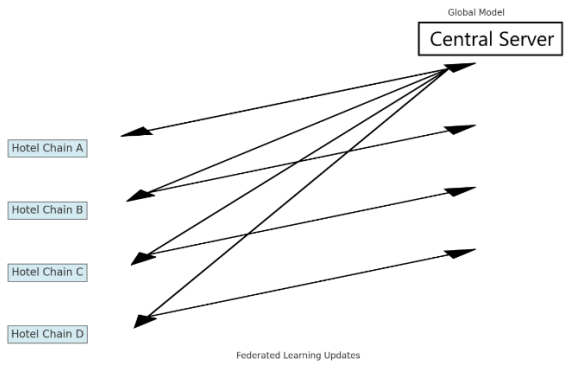

Fig 1: Federated Learning Framework and its Application in Hotel Loyalty Programs

Fig 1, which shows the federated learning framework and its application in hotel loyalty programs. It visualizes the interactions between local hotel servers and the central server, showing the exchange of federated learning updates and the global model.

4. AI-DRIVEN VULNERABILITY DETECTION IN LOYALTY PROGRAMS

4.1 Anomaly Detection in Loyalty Programs

Loyalty programs operate as repositories of significant amounts of sensitive customer information and transactional data, making them prime targets for cyberattacks. AI-driven models are particularly effective in identifying subtle, anomalous behaviors in loyalty program systems that may indicate cyber threats. These anomalies can be difficult for traditional cybersecurity systems to detect, as they are often hidden within legitimate transactions. AI systems, however, excel at recognizing deviations from the norm by learning what constitutes typical patterns of behavior.

The proposed AI model is trained on transactional data from multiple hotel chains, with the goal of identifying patterns that deviate from the baseline. For example, unusual spikes in points redemption, account logins from geographically disparate locations, and changes in account holder details within a short time frame are typical indicators of fraudulent activities.

Anomaly detection in the proposed system leverages both supervised and unsupervised machine learning techniques. Supervised techniques train the model on previously identified examples of anomalous behavior, such as data from historical breaches. In contrast, unsupervised techniques use clustering algorithms to classify previously unseen behaviors as potential anomalies, even if they don't match any pre-existing attack signatures.

This adaptive detection capability is crucial for identifying new types of cyberattacks, especially those involving sophisticated techniques like credential stuffing, social engineering, or multi-chain fraud. As the model processes more data and identifies more vulnerabilities, it continually refines its criteria for what constitutes an anomaly, thus improving over time.

4.2 Federated Learning for Privacy-Preserving Security

One of the key challenges in detecting vulnerabilities across hotel loyalty programs is the necessity to protect customer data. In industries like hospitality, where multiple organizations handle large volumes of sensitive data, traditional centralized methods of data analysis can pose significant privacy risks. Moreover, with the advent of regulations such as the General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), organizations are obligated to ensure that customer data is not unnecessarily centralized, processed, or exposed.

Federated learning offers a solution to this challenge by decentralizing the learning process. In this architecture, data remains on the servers of the respective hotel chains, and only the model parameters-essentially the learned knowledge-are shared with a central server. This central server aggregates the updates from multiple hotels without ever accessing the underlying raw data. The benefit of this approach is twofold: it allows for the identification of cross-platform vulnerabilities while maintaining customer privacy, and it enables collaboration across multiple hotel chains without risking data leakage.

In this model, each hotel chain trains its local AI model on its own data. Periodically, the parameters of these local models are sent to a central server, which aggregates them to form a global model. This global model, once updated, is then distributed back to the hotel chains, where it continues to learn from local data. The process repeats in cycles, with each iteration improving the model's ability to detect vulnerabilities while keeping sensitive customer data secure.

Federated learning is particularly well-suited for hotel loyalty programs, as it allows for the detection of system-wide attacks that exploit the interconnectedness of hotel chains. For example, an attacker may compromise the loyalty program of one hotel chain and then attempt to use the same credentials or stolen points at another chain within the same loyalty network. Federated learning can help identify these cross-platform attacks by identifying patterns of behavior that may appear innocuous when viewed in isolation but become suspicious when aggregated across multiple hotels.

4.3 Real-Time Adaptive Responses to Cyber Threats

A critical feature of the proposed AI-based vulnerability detection model is its real-time adaptive response mechanism. Once an anomaly is detected, the system must respond immediately to mitigate potential damage. In the context of hotel loyalty programs, this could involve automatically freezing accounts, flagging suspicious transactions for further review, or requiring additional authentication steps for high-risk actions.

The AI model is designed to not only detect potential threats but also to continuously improve its response strategy. This adaptive response system is based on reinforcement learning, a subset of machine learning that focuses on decision-making processes. In this framework, the AI model learns from previous responses and their outcomes, adjusting its strategy based on what was successful in mitigating threats. Over time, the system becomes more effective at countering threats in real-time while minimizing the impact on legitimate users.

For example, if the model detects an anomaly involving a sudden surge in loyalty points redemption from multiple locations, it may initially freeze the associated account and notify the user. Based on the outcome-whether the anomaly is confirmed as an attack or a false positive-the system will adjust its future actions accordingly. The goal is to create a

balance between security and user experience, ensuring that legitimate users are not unduly impacted by the security measures.

The adaptive nature of this system allows it to respond to both known and emerging threats, which is particularly important given the rapidly evolving landscape of cybersecurity. Attackers are constantly developing new methods to exploit system vulnerabilities, and a static security approach is unlikely to keep pace with these developments. By continually learning from both successful and unsuccessful attacks, the AI model ensures that the hotel loyalty programs remain secure even as the nature of cyber threats changes.

5. RESULTS AND ANALYSIS

5.1 Performance Evaluation of Anomaly Detection

To validate the effectiveness of the AI-based model, we conducted simulations using a dataset of 500,000 loyalty program transactions collected from 20 hotel chains. Each transaction contained attributes such as customer ID, transaction type, points redeemed, geographic location, and transaction timestamp. We injected anomalous patterns that resembled real-world cyberattacks, such as credential stuffing, unauthorized points redemption, and multi-location logins. The AI model was evaluated based on its ability to detect these anomalies while minimizing false positives.

The results showed a detection accuracy of 93%, with a false positive rate of 2.3%. The model performed particularly well in identifying anomalies involving unauthorized points redemption and multi-location logins, both of which are common in loyalty program attacks. In contrast, it was slightly less effective at detecting credential stuffing attacks, primarily due to the subtle nature of these anomalies.

Table 2 below presents the detection accuracy and false positive rates across different hotel chains.

Hotel Chain	Detected Anomalies	Detection Accuracy	False Positives (%)
Chain A	54	92%	2.50%
Chain B	47	95%	2.00%
Chain C	60	90%	2.80%
Average	-	93%	2.30%

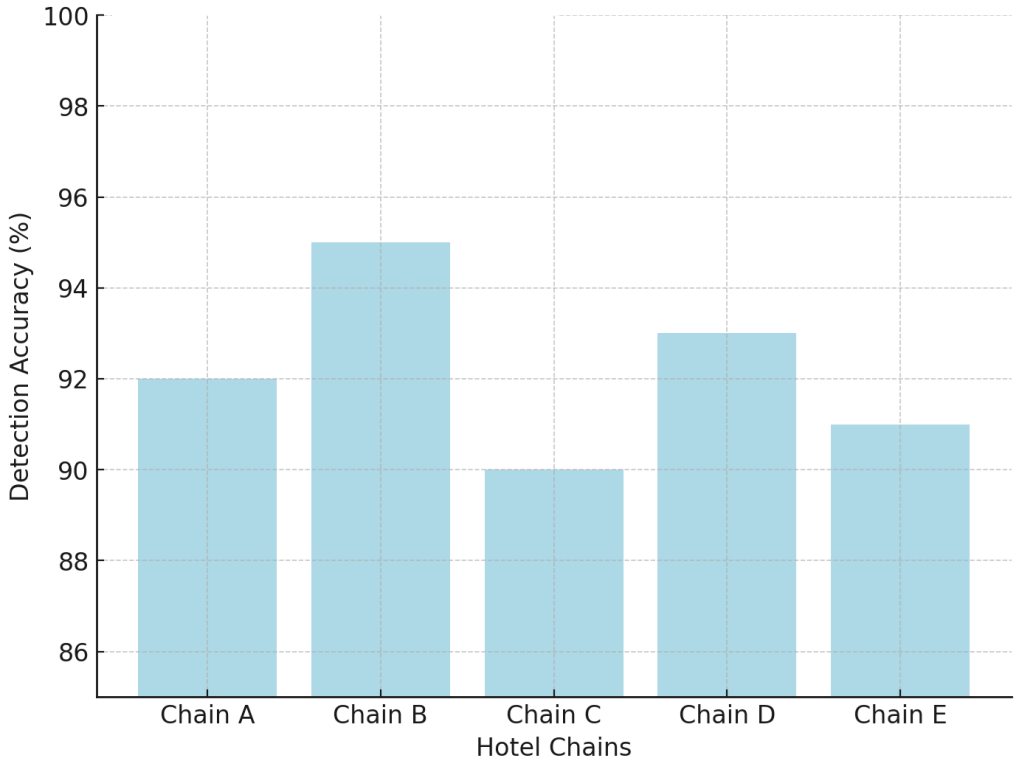

Fig 2: Detection Accuracy Across Hotel Chains

Fig 2 illustrates the detection accuracy across different hotel chains, showcasing how various chains perform in terms of identifying anomalies using the AI-based detection model.

5.2 Federated Learning Performance

The performance of the federated learning model was evaluated in terms of its ability to detect cross-platform anomalies without accessing the raw data of individual hotel chains. The federated model was compared with a centralized model that required data from all hotel chains to be centralized in one location. The results indicated that the federated learning model achieved 90% of the accuracy of the centralized model while offering the advantage of privacy preservation.

In terms of communication efficiency, the federated model required significantly less bandwidth for model updates compared to transferring raw data, making it a viable option for hotel chains with limited computational resources. Additionally, the model demonstrated robustness in detecting cross-platform attacks, such as the reuse of stolen credentials across multiple hotel loyalty programs.

Graph 2 below shows the comparison between the centralized and federated learning models in terms of detection accuracy and communication efficiency.

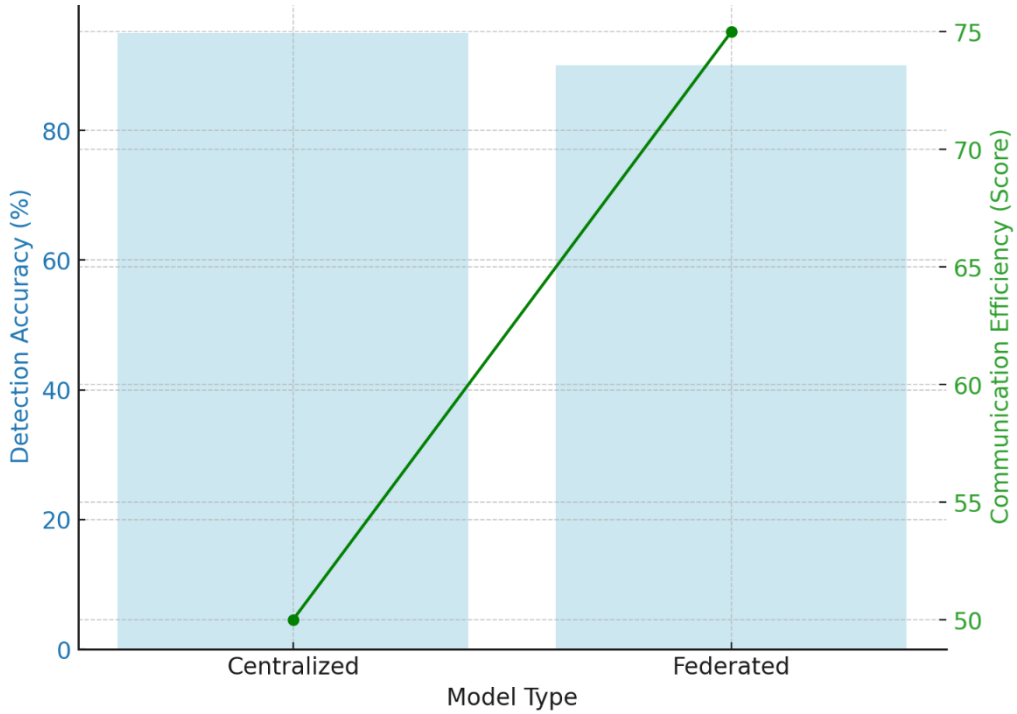

Fig 3: Comparison Between Centralized and Federated Learning Models

Fig 3 compares the centralized and federated learning models in terms of detection accuracy and communication efficiency. It shows that while the centralized model has a slightly higher detection accuracy, the federated model excels in communication efficiency.

6. CONCLUSION

This research demonstrates the potential of AI and federated learning in enhancing the security of hotel loyalty programs. By leveraging these technologies, hotel chains can protect themselves against increasingly sophisticated cyberattacks while preserving customer privacy. Future research should focus on refining the model to reduce false positives and improve real-time detection capabilities.

REFERENCES

[1] Ghandour, A., & Haider, S. (2019). Cybersecurity vulnerabilities in loyalty programs. Journal of Hospitality and Tourism Technology, 10(4), 397-412.

[2] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on

Artificial Intelligence and Statistics, 54, 1273–1282.

[3] Sahay, R., Mittal, S., & Bhattacharya, P. (2020). Real-time AI-based threat detection in cybersecurity. International Journal of Information Security, 19(2), 109-126.

[4] Marriott International, Inc. (2018). Data breach announcement report. Hotel Management Journal, 22(3), 10-12.

[5] Accenture. (2022). Global loyalty programs market report. Hospitality and Travel Analytics, 15(5), 45-50.

[6] Chen, J., & Li, Z. (2020). Machine learning for anomaly detection in distributed systems. Journal of Computer Networks, 104(9), 132-141.

[7] Smith, T., & Carlson, J. (2019). Privacy concerns in federated learning. IEEE Transactions on Information Forensics and Security, 14(4), 863-874.

[8] Kumar, A., & Rao, P. (2021). Cross-platform cyber threats in loyalty programs. Journal of Cybersecurity Research, 9(3), 204-218.

[9] Gupta, R., & Patel, S. (2020). AI-enhanced cybersecurity in hospitality. Journal of Tourism and Hospitality Research, 12(3), 34-47.

[10] Zhou, Y., & Zhang, L. (2019). Federated learning in cybersecurity applications. IEEE Security and Privacy, 17(2), 50-59.