Dynamic Denial of Service Attack Prevention using a Multi-Level IP Shedding Defense Mechanism

By

Richard Wambua Kiundi

153031

Submitted in partial fulfillment of the requirements for the Degree of Master of Science in Information Systems Security at Strathmore University

School of Computing & Engineering Sciences

Strathmore University

Nairobi, Kenya

June, 2025

Declaration and Approval

Declaration

I declare that this work has not been previously submitted and approved for the award of a degree by this or any other University. To the best of my knowledge and belief, the dissertation contains no material previously published or written by another person except where due reference is made in the dissertation itself.

Richard Wambua Kiundi

Signature ……

19th May 2025

Date…….

Approval

The dissertation of Richard Wambua Kiundi was reviewed and approved by the following:

Dr. Vitalis Ozianyi

Lecturer, School of Computing & Engineering Sciences,

Strathmore University

Dr. Julius Butime,

Dean, School of Computing & Engineering Sciences,

Strathmore University

Prof. Bernard Shibwabo,

Director of Graduate Studies,

Strathmore University

Abstract

Distributed Denial of Service (DDoS) attacks continue to evolve in scale and sophistication, overwhelming traditional defenses that struggle to distinguish malicious traffic from legitimate users. This dissertation proposes a novel Multi-Level IP Shedding Defense Mechanism (MLISDM) to dynamically mitigate DDoS threats through real-time IP reputation analysis and granular geographic zoning. By correlating IP addresses with localized zones (e.g., towns, regions) via WHOIS database insights, the system enforces adaptive, tiered restrictions (Levels 0–6) to block malicious traffic while minimizing disruption to legitimate users.

The framework integrates a behavior-based Intrusion Detection System (IDS) and Prevention System (IPS), enabling real-time traffic analysis and firewall rule adjustments during attacks. Unlike static Access Control Lists (ACLs), this approach reduces collateral damage by leveraging geographic intelligence and IP reputation scoring to prioritize high-risk zones. Prototype testing on a private network testbed demonstrated the system's ability to neutralize simulated DDoS attacks with over 90% accuracy while maintaining service availability for legitimate users.

Key innovations include the novel integration of profitability metrics with granular user location data. This synergy refines defense precision by enabling business-value-driven traffic prioritization—a capability often absent in conventional, technically-focused DDoS mitigation techniques—which in turn ensures critical assets receive prioritized protection during attacks. The results highlight MLISDM's scalability, adaptability, and reduced reliance on manual intervention, offering a proactive defense against evolving DDoS tactics. This work advances information security by bridging gaps in current mitigation strategies, providing a blueprint for intelligent, context-aware cybersecurity frameworks.

Keywords: DDoS mitigation, IP reputation, geographic zoning, adaptive access control, intrusion prevention.

Table of Contents

Declaration and Approval ii
Abstract iii
Table of Contents iv
List of Figures viii
List of Tables ix
Acknowledgements x
Dedication xi
Chapter 1: Introduction 1
1.1 Background to the Study 1
1.2 Problem Statement 2
1.3 Research Objectives 3
1.4 Research Questions 3
1.5 Scope of the Research 3
1.6 Limitations of the Research 4
1.7 Justification of the Research 4
Chapter Two: Literature Review 5
2.1 Introduction 5
2.2 Technology Review of DDoS Attacks 5
2.2.1 Types of DDoS Attacks and Targeted Services 5
2.2.2 Technology Vulnerabilities Exploited by DDoS Attacks 7
2.3 Theoretical Review of DDoS Attacks 8
2.4 Empirical Review of DDoS Attacks 8
2.5 Intrusion Detection Systems 10
2.6 Intrusion Prevention Systems 11
2.7 DDoS Mitigation Strategies 11
2.7.1 DDoS Mitigation Tools 11
2.7.2 DDoS Mitigation Systems 12

2.7.3 Use of IP Address Reputation in Intrusion Mitigation 12
2.8 Global Coordination of IP Addresses and AS Numbers by IANA 13
2.9 Research Gaps 14
2.10 Conceptual Framework 15
Chapter Three: Research Methodology 18
3.1 Introduction 18
3.2 Research Design 18
3.3 System Development Methodology 19
3.3.1 Requirements Specification 19
3.3.2 System Design 20
3.3.3 System Analysis 20
3.3.4 System Development 20
3.3.5 System Testing 20
3.4 System Validation 21
3.5 Ethical Considerations 22
CHAPTER 4: System Design and Architecture 24
4.1 System Overview 24
4.2 Overall System Architecture 25
4.3 IP Shedding Algorithm 28
4.4 Data Flow and Interactions 29
Chapter 5: Implementation And Testing 33
5.1 Introduction 33
5.2 System Environment 33
5.3 Testbed Setup 33
5.4 IP Address Location Classifier Component 35
5.5 Profit Classifier Component 38
5.5.1 Conceptual Data Sources and Integration for Profit Classification 38
5.5.2 Data Scanned/Used 39
5.5.3 Prototype Implementation (Simplified Emulation): 40

5.6 DDoS Integrated Defense Mechanism 41
5.7 Testing of Defense Mechanism 44
5.7.1 Server Load Measurement Tool 44
5.7.2 Slowloris DDoS Attack Tool 44
5.7.3 Client Request Tool 46
5.7.4 Client Baseline Tests 46
5.7.5 DDoS Performance Tests 49
5.8 Revenue Preservation 54
Chapter 6: Discussion of Results 56
6.1 Introduction 56
6.2 Effectiveness of the Multi-Level IP Shedding Approach 56
6.3 Performance Metrics Analysis 57
6.3.1 Server Load Management 57
6.3.2 Client Performance Preservation 57
6.4 Revenue Preservation Analysis 58
6.5 Comparison with Existing DDoS Defense Mechanisms 58
6.5.1 Core Intelligence Integration in MLISDM 59
6.5.2 Comparative Advantages of the MLISDM Approach 60
6.6 Limitations and Challenges 61
6.6.1 Computational Resource Constraints 61
6.6.2 Traffic Classification Challenges 62
6.6.3 Scalability Considerations 62
6.7 Implications for DDoS Defense Strategies 62
6.7.1 Business-Aligned Security Posture 62
6.7.2 Geographical Traffic Prioritization 62
6.7.3 Dynamic Defense Adaptation 63
6.8 Summary of Findings 63
Chapter 7: Conclusions, Recommendations and Future Work 64
7.1 Conclusions 64

7.2 Research Contributions 65
7.3 Recommendations 65
7.4 Limitations of the Study 66
7.5 Future Work 67
7.6 Overall Research Assessment and Reflections 68
References 71
Appendices 76
Appendix A: Similarity Report 76
Appendix B: Ethical Clearance Confirmation 77
Appendix C: PHP code for IP Location Classifier Component 78
Appendix D: PHP code for Profit Classifier Component 80
Appendix E: Bash Script To Measure Server Load 82
Appendix F: PHP Script for Client Request Tool 83
Appendix G: Web server application code 84
Appendix H: Bash script to activate iptables rules 85

List of Figures

Figure 2.1: Location and Probability Classifier……...17
Figure 4.1: System Architecture Diagram.…..……..26
Figure 4.2: Sequence Diagram - Attack Detection and Mitigation… ……29
Figure 4.3: Data Flow Diagram……...……..32
Figure 5.1: IP location classifier component… ……..…….36
Figure 5.2: IP zone details table ……...…….……..37
Figure 5.3: Profit Classifier component……40
Figure 5.4: DDoS Integrated Defense Mechanism……...42
Figure 5.5: Web server performance for Slowloris baseline test……..45
Figure 5.6: Server load with one client – 100 requests every 5 seconds……..46
Figure 5.7: Client Performance – One client at 100 request every 5 seconds ……..47
Figure 5.8: Server load – two clients with 100 requests……47
Figure 5.9: Client performance – Two client with 100 requests……48
Figure 5.10: Client performance - Two clients with 200 requests……..49
Figure 5.11 Server load under unmitigated DDoS attack……...50
Figure 5.12: Client performance during DDoS attack……..50
Figure 5.13: Server load under mitigated DDoS attack……51
Figure 5:14: Performance of client 2 under a mitigated attack……..52
Figure 5:15: Performance of client 1 under a mitigated attack……..52
Figure 5.16: Revenue Preservation Comparison……...55

List of Tables

Table 5.1: Profitability data……...54

Acknowledgements

A special mention to my supervisor Dr. Vitalis Ozianyi, together with Dr. Joseph Sevilla, and the Strathmore University staff especially those at @iLab Africa and the Faculty of Information Technology for the great support I got from them and good relations throughout the time I've been a student at this great university.

Dedication

I dedicate this dissertation to my sons Nzalu and Mumo. You are an inspiration and a joy to have in my life.

Chapter 1: Introduction

1.1 Background to the Study

The rise of the Internet has significantly altered the landscape of communication, business operations, and information access. Yet, this digital advancement has also brought about a host of security challenges, with Distributed Denial of Service (DDoS) attacks standing out as especially harmful. DDoS attacks are a specific type of cyber assault aimed at making computer resources, like websites or networks, inaccessible to legitimate users by overwhelming them with excessive online traffic. These attacks exploit the inherent features of the Internet, turning its interconnectedness and openness into exploitable vulnerabilities (Kumar, Dwivedi, Kumar, & Gill, 2024). Over time, DDoS attacks have evolved from minor nuisances to intricate, highly disruptive events that can affect organizations of various sizes across different sectors (Dupont, Shearing, Bernier, & Leukfeldt, 2023). The motivations driving these attacks are diverse, including financial incentives and political aims, while the tactics used continue to advance in sophistication (Shafiq, Rizwan, & Alhassan, 2022).

Traditional countermeasures, such as firewalls and Intrusion Detection Systems (IDS), have served as the primary line of defense against DDoS incidents. However, these conventional techniques often fall short when faced with the scale and intricacy of modern DDoS attacks (Douglas, Heath, & Thompson, 2020). A prominent challenge in mitigating such attacks is the difficulty in differentiating between legitimate and malicious traffic. Attackers often utilize spoofed IP addresses, complicating the identification of the attack's source. Furthermore, the sheer volume of incoming traffic can easily overwhelm standard security systems, leading to significant downtime and substantial financial losses for the targeted networks (Mallick & Nath, 2024). This highlights an urgent need for more adaptive and dynamic defense systems capable of reacting in real-time to the constantly changing threat landscape.

The concept of IP reputation has emerged as a valuable technique for bolstering defenses against DDoS attacks. By evaluating the historical involvement of IP addresses in harmful activities, security systems can make more informed decisions about which traffic to permit or block (Husák & Záhonová, 2020). The WHOIS database (Internet Corporation for Assigned Names and Numbers, n.d) plays a crucial role in this method, as it provides essential information about the registered users of IP addresses, offering key insights into their geographic origins and relevant contact details.

Building on the concept of IP reputation, this dissertation proposes a multi-tiered IP shedding defense mechanism that employs the WHOIS database to enforce zone-based access restrictions. By linking IP addresses to smaller geographic areas, this system can implement targeted restrictions aimed at minimizing disruptions for legitimate users while effectively addressing potential threats (Imthiyas & Handan, 2020). This sophisticated approach allows for a more refined level of control, which is vital in an environment where cyber attackers can quickly adjust their tactics and alter their IP addresses.

The proposed defense mechanism depends on the integration of a behavior-based Intrusion Detection System (IDS) alongside an Intrusion Prevention System (IPS). The IDS monitors network traffic for indications of a DDoS attack, such as unexpected spikes in traffic or deviations from normal behavior patterns (Pasupathi, Kumar, & Pavithra, 2025). Once anomalies are identified, the IPS dynamically modifies firewall rules to implement the appropriate level of IP shedding based on the zone's reputation. This adaptive strategy is essential for maintaining continuous service availability while minimizing the attack's impact.

Further enhancement of this defensive mechanism can be achieved by incorporating insights related to user location and profitability. Location-aware DDoS protection utilizes the distributed nature of attacks to allow defenders to analyze traffic origins and evaluate their alignment with the expected geographic profiles of legitimate users (Zilberman, et al., 2024). This strategy has the potential to identify and mitigate traffic from unexpected locations, making it especially beneficial for businesses operating within specific regions or countries. Moreover, it can protect online applications crucial to Content Distribution Networks (CDNs) that function within defined geographic areas.

1.2 Problem Statement

The swift advancement and widespread occurrence of Distributed Denial of Service (DDoS) attacks present a significant threat to the integrity and security of Internet-based systems. Despite the progress made in traditional security solutions such as firewalls and Intrusion Detection Systems (IDS), these strategies often prove inadequate in effectively countering the scale, complexity, and adaptability of contemporary DDoS operations. Attackers take advantage of fundamental vulnerabilities within the Internet's architecture, utilizing techniques such as IP spoofing and randomized source IP addresses to inundate target servers with fraudulent traffic. This situation complicates the ability of conventional defense mechanisms to differentiate between legitimate users and malicious entities, resulting in considerable service interruptions and financial repercussions.

Current methodologies frequently lack the accuracy and responsiveness necessary to adapt to swiftly evolving attack patterns. Static Access Control Lists (ACLs) and standard traffic blocking strategies often lead to significant collateral damage, inadvertently restricting access for legitimate users and diminishing the effectiveness of these systems. Moreover, the lack of a comprehensive mechanism that integrates IP reputation, geographic intelligence, and dynamic response capabilities has created a substantial void in DDoS defense frameworks.

1.3 Research Objectives

i. To explore the characteristics, consequences, and strategies associated with DDoS attacks.
ii. To assess the current technologies available for DDoS mitigation.
iii. To conceptualize, create, and evaluate a multi-tier IP shedding defense strategy.
iv. To analyze the efficacy of the multi-tier IP shedding defense mechanism against DDoS attacks.

1.4 Research Questions

i. What are the predominant patterns observed in Distributed Denial-of-Service (DDoS) attacks, their repercussions, and the strategies employed in recent incidents?
ii. How do current DDoS mitigation technologies measure up in terms of efficacy and their inherent limitations?
iii. What architectural and design factors should be considered when developing a robust multi-level IP shedding defense mechanism for DDoS mitigation?
iv. In what ways does the multi-level IP shedding defense mechanism differ from conventional DDoS mitigation strategies regarding accuracy, performance implications, and the speed of threat neutralization?

1.5 Scope of the Research

This study is designed to develop and assess an innovative multi-tier IP shedding defense mechanism aimed at alleviating the growing threat posed by Distributed Denial of Service (DDoS) attacks. The proposed defense strategy will undergo thorough evaluation within a controlled private network testbed, where it will simulate realistic DDoS attacks targeting internet-based services. Performance metrics will be utilized to evaluate effectiveness, the impact on legitimate users, and the speed of mitigation, with comparisons made to conventional defense mechanisms. This research will not address the security of physical infrastructure, non-DDoS cyber threats, or attacks specifically directed at mobile devices. The focus of the analysis will remain primarily on IP-based DDoS attacks.

1.6 Limitations of the Research

Recognizing the limitations of this study is essential for comprehending the context and applicability of its findings. While the research seeks to offer innovative strategies for mitigating DDoS attacks, it is conducted within specific constraints that warrant acknowledgment. Firstly, the use of a controlled testbed is crucial for the development and evaluation of the proposed solutions; however, it inherently reduces the complexity found in real-world DDoS attacks and the ever-evolving dynamics of the internet. Secondly, concentrating on specific geographical areas may yield valuable localized insights but may also restrict the generalizability of the mechanism to the broader spectrum of global DDoS tactics. Furthermore, although the study examines a wide array of attacks, it cannot predict every conceivable future method of attack. The defense mechanism proposed will depend on existing technologies related to IP reputation, zone mapping, and intrusion detection, which may require updates as these technologies progress.

1.7 Justification of the Research

The increasing severity and sophistication of Distributed Denial of Service (DDoS) attacks necessitate the advancement of effective defense mechanisms to safeguard the availability and integrity of internet-based services. Conventional firewalls and intrusion detection systems frequently struggle to accurately differentiate between legitimate and malicious traffic, particularly when assailants employ spoofed IP addresses and adapt their attack strategies. This inadequacy renders networks susceptible to significant downtime and the potential for recurrent attacks that can exacerbate the overall damage. As DDoS attacks become more frequent, intricate, and accessible—largely due to the proliferation of botnets-for-hire and DDoS-for-hire services, the pressing need for dynamic and adaptive defense solutions to address these evolving threats become evident. This research introduces an innovative multi-level IP shedding defense mechanism that utilizes real-time IP address zone reputation along with behavior-based intrusion detection to enhance the precision and proactivity of DDoS mitigation efforts. Internet applications characterized by a predominant user base originating from a specific geographical region stand to gain significantly from the proposed DDoS solution. These include online banking platforms, mobile-money merchant checkouts, and content delivery networks (CDNs).

Chapter Two: Literature Review

2.1 Introduction

This chapter provides a comprehensive review of the literature pertaining to Distributed Denial of Service (DDoS) attacks, strategies for mitigation, and notable recent instances of such attacks. It additionally examines the global distribution of IP addresses, and the methodologies employed to ascertain the physical location of a device based on its IP address. The chapter concludes with the presentation of a conceptual framework for the proposed DDoS mitigation solution.

2.2 Technology Review of DDoS Attacks

"Illegitimate user traffic" constitutes a major and serious cyber threat, particularly exemplified by Distributed Denial of Service (DDoS) attacks, which aim to disrupt the availability of network resources. Unlike traditional Denial of Service (DoS) attacks that stem from a single source, DDoS attacks leverage multiple compromised devices, often referred to as botnets, to bombard a target site with overwhelming traffic, thus making services unavailable to legitimate users (Gupta & Dahiya, 2021). A prominent case is the Dyn attack in October 2016, where the Mirai botnet infiltrated thousands of Internet of Things (IoT) devices, resulting in extensive outages for major platforms such as Twitter, Netflix, and GitHub (Mohammed, Eze, & Alhassan, 2023). The intricate nature of distributed attacks complicates their detection and prevention efforts (Kumar, Sharma, & Singh, 2020). The underlying motives for DDoS attacks can vary significantly, including acts of sabotage, political objectives, or the intent to disrupt services (Shafiq, Rizwan, & Alhassan, 2022). Frequently targeted sectors encompass e-commerce, financial services, businesses, and governmental organizations, where service interruptions can cause considerable financial repercussions, damage to reputation, or lowered operational efficiency (George, Baskar, & Srikaanth, 2024). Another significant incident occurred in 2018 when GitHub experienced an attack that peaked at an astonishing 1.35 Tbps, utilizing Memcached amplification techniques (Gupta & Singh, 2020). Attackers regularly exploit inadequately secured IoT devices and other vulnerable internet systems to carry out DDoS attacks (Esquivel, Llewellyn, & Khan, 2020).

2.2.1 Types of DDoS Attacks and Targeted Services

DDoS attacks are generally classified into three main types: volumetric, protocol-based, and application-layer attacks. Volumetric attacks seek to create congestion within the targeted network, resulting in bandwidth depletion. Attackers often employ fundamental protocols, such

as DNS amplification or UDP floods, where small requests generate significantly larger responses aimed at the target (Mansoor, Malik, & Khan, 2021). These types of attacks are relatively easy to execute using botnets and can severely affect internet-facing services (Nour, Murtaza, & Rao, 2022). Protocol-based attacks exploit vulnerabilities in network communication protocols, often targeting weaknesses in how systems manage or interpret specific protocols. A common example is the SYN flood attack, which takes advantage of the TCP handshake process. In this scenario, attackers bombard a server with SYN packets to initiate connection requests but fail to send the final ACK packet, resulting in the server holding a backlog of half-open connections and ultimately depleting resources needed for legitimate requests (Sajjad & Khan, 2020). Additionally, protocol-based attacks can involve the exploitation of packet fragmentation during data transfers. Maliciously fragmented packets can be created, which, when reassembled by the targeted system, can exploit vulnerabilities in the handling of fragmented data. The historical Teardrop attack is a notable example of this, having successfully exploited weaknesses in the reassembly of overlapping IP fragments in older Windows systems. Such attacks complicate defense measures since they may closely resemble legitimate traffic. Traditional firewalls, which primarily evaluate only packet headers, might not effectively detect these threats. Conversely, Deep Packet Inspection (DPI) offers a more sophisticated approach that examines data packet contents to identify malicious patterns or anomalies (Hassan, Javed, & Baig, 2022). While DPI can successfully detect and mitigate threats like SYN floods and fragmentation exploits, it can also raise privacy concerns due to its capability to analyze user data (Yang & Alhassan, 2022).

Application-layer attacks are deemed the most complex variety of DDoS attacks because they often appear legitimate in their execution. These attacks usually target specific applications, such as web servers and databases, with the intent of exhausting application resources. HTTP floods typify this attack type, where attackers aim to inundate the server with a large volume of HTTP requests (Khan & Khan, 2021). Most DDoS attacks occur at the application layer, as the requests generally disguise themselves as legitimate; therefore, services require advanced Intrusion Detection Systems (IDS) that apply behavioral analysis techniques to differentiate between legitimate users and malicious connections.

DNS application-layer DDoS attacks are primarily conducted by overwhelming DNS servers with excessive requests, leading to significant service interruptions. Attackers frequently use techniques such as DNS amplification, exploiting the intrinsic characteristics of the DNS protocol. In this technique, a small request is sent to a DNS server, which subsequently

produces a much larger response directed at the target's IP address (Sullivan, 2018). This amplification effect significantly increases the traffic directed at the target, overwhelming its resources. Alternatively, attackers may deploy botnets—networks of compromised devices that collectively execute an assault. By leveraging the aggregated bandwidth of multiple infected machines, a botnet can generate substantial volumes of DNS queries aimed at the DNS server, causing it to become unresponsive and ultimately disrupting access to services dependent on that server (Alhassan, Eze, & Mohammed, 2023). The 2016 Dyn DNS cyber-attack serves as a pertinent illustration of this attack type. This incident involved a significant surge of DNS requests targeting Dyn, a leading DNS provider. While it is acknowledged that various attack vectors were employed, the result was a marked degradation of services for popular platforms like Twitter, Netflix, and Spotify. This attack exposed the vulnerabilities of DNS systems to both amplification assaults and botnet-driven traffic, resulting in widespread service disruptions across multiple sectors (Sullivan, 2018). The event underscored the pressing need for enhanced security measures to protect DNS infrastructure against such coordinated and complex DDoS attacks.

2.2.2 Technology Vulnerabilities Exploited by DDoS Attacks

In the realm of Distributed Denial of Service (DDoS) attacks, it is crucial to analyze various technological vulnerabilities that may lead to such threats. These include DNS amplification attacks, shortcomings in Internet of Things (IoT) devices, and weaknesses in cloud relay services. A significant vulnerability lies in DNS amplification attacks, where cybercriminals exploit publicly accessible DNS resolvers to amplify attack traffic. This method relies on the ability of a small DNS query to generate a disproportionately large response directed at the target, resulting in an overwhelming influx of traffic that can incapacitate its resources (Davis et al., 2020). By using public DNS services, users inadvertently expose themselves to security risks, becoming susceptible to exploitation without direct attacks on these servers (Alhassan, Eze, & Mohammed, 2023).

Another critical vulnerability is found in IoT devices. These devices often come with default settings or lack adequate security measures, making them prime targets for malicious actors. Once compromised, these devices can be incorporated into botnets that facilitate large-scale DDoS attacks. A notable example of this exploitation is the Mirai botnet, which conducted extensive attacks using numerous unsecured IoT devices (Fahim, Alhassan, & Mohammed, 2021). This situation highlights a significant security gap that necessitates urgent attention to protect these pervasive technologies (Lee & Min, 2022).

Moreover, DDoS attackers seize upon vulnerabilities inherent in cloud relay services. As organizations continue to shift their operations to the cloud, the threat of DDoS attacks intensifies. Poorly configured cloud services or insecure security protocols can provide attackers with opportunities to launch denial-of-service campaigns. The distributed nature of cloud delivery models can further exacerbate the threat landscape for applications and services hosted in the cloud (Kumar, Singh, & Srivastava, Vulnerabilities in cloud services and the implications for DDoS attacks, 2022). Identifying these vulnerabilities is essential, as it enables the development of effective preventative measures, distinguishing between persistent weaknesses and emerging risks that may arise in the continuously evolving digital environment (Nour & Murtaza, 2020).

2.3 Theoretical Review of DDoS Attacks

The theoretical foundation for understanding Distributed Denial-of-Service (DDoS) attacks is rooted in Game Theory. This framework facilitates the analysis of interactions between attackers and defenders, as outlined by Nash, von Neumann, and Morgenstern in their seminal works. Game Theory posits that all participants act rationally to achieve optimal outcomes (Ho, Rajagopalan, Skvortsov, Arulampalam, & Piraveenan, 2022) Within the realm of DDoS attacks, the objective of attackers is to maximize disruption to the availability of services, while defenders strive to implement measures that mitigate such disruptions and preserve accessibility (Bert, Zhang, & Luo, 2020). Utilizing Game Theory allows researchers to develop robust defensive strategies by evaluating the probabilities of losing both legitimate and malicious users (Li, Liu, & Zhang, 2022). This analytical framework deepens the understanding of the behaviors and strategies that characterize DDoS attacks, thereby equipping network administrators with the insights needed to create more effective countermeasures (Li, Liu, & Zhang, 2022). Moreover, the application of best response models enables defenders to decode the tactics employed by attackers, ensuring that networks remain functional and protected from potential threats (Zhang & Liu, 2020).

2.4 Empirical Review of DDoS Attacks

One of the most notable recent Distributed Denial of Service (DDoS) attacks was the Dyn DNS attack, which transpired in October 2016. Executed via the Mirai botnet, this assault compromised hundreds of thousands of Internet of Things (IoT) devices around the world. Devices such as poorly secured security cameras and home routers were repurposed to inundate the Dyn DNS infrastructure, generating an overwhelming amount of traffic that disrupted

various internet services. Websites such as Twitter, Netflix, and GitHub were among those significantly affected, demonstrating the far-reaching impacts that a single attack can have (Mansoor, Malik, & Khan, 2021). This incident highlights the critical need to secure IoT devices and other interconnected systems to prevent their exploitation in large-scale attacks (Kumar, Sharma, & Singh, 2020).

Another significant event occurred in 2012 with an attack on Cloudflare, where the attack volume soared to an unprecedented level exceeding 300 Gbps. This incident utilized both DNS and NTP amplification techniques to reflect and increase the traffic directed at the target, showcasing how vulnerabilities in public services and infrastructure can be exploited to generate considerable traffic even against robust defenses like those provided by Cloudflare (Alhassan, Eze, & Mohammed, 2023). The attack led to a reevaluation of mitigation strategies by companies, driving advancements in DDoS protection measures.

In 2018, the attack on GitHub illustrated the evolving nature of DDoS attacks, marking it as one of the largest incidents recorded with traffic peaking at 1.35 Tbps. This assault employed Memcached amplification, which exploits unsecured Memcached instances to overwhelm the target server. Although GitHub utilized Akamai's DDoS mitigation services to maintain high availability, the attack underscored the need for ongoing improvement in defenses against both current and emerging DDoS threats (Gupta & Singh, 2020).

Moreover, the Anonymous Sudan attacks that targeted Kenya's eCitizen platform and Safaricom in July 2023 provide a compelling case for understanding the changing landscape of DDoS attacks and their potential ramifications. eCitizen acts as the government’s primary portal for accessing over 5,000 public services, while Safaricom serves as the leading telecommunications provider crucial for the M-PESA mobile money service. These platforms faced significant disruptions due to a series of attacks orchestrated by the Anonymous Sudan hacktivist group, suspected to have connections to pro-Russian entities. Employing advanced layer 7 DDoS attack strategies, this group flooded these services with malicious traffic aimed at disrupting specific web applications (Kumari & Jain, 2023). The ramifications were extensive, hindering access to vital government services, obstructing financial transactions, and affecting mobile communications for millions of Kenyan citizens (Mohammed, Eze, & Alhassan, 2023).

The repercussions were profound, as reliance on these digital platforms not only impacted businesses and government operations but also the everyday lives of individuals. Despite efforts from the Kenyan government and Safaricom to mitigate the attacks and provide

alternative access options, the incident exposed significant vulnerabilities in the nation's digital infrastructure (Alhassan, Eze, & Mohammed, 2023). This situation accentuated the pressing requirement for comprehensive cybersecurity measures, which include advanced threat detection systems, proactive incident response plans, and continuous investments to enhance digital defenses (Esquivel, Llewellyn, & Khan, 2020). Additionally, these attacks highlighted the critical importance of international collaboration in addressing cyber threats that transcend national borders. A detailed examination of the technical specifics of the attacks, the response strategies employed, and the overall societal impact can offer invaluable insights for governments and organizations globally as they work to bolster their cybersecurity resilience against increasingly sophisticated and disruptive cyber threats (Shafiq, Rizwan, & Alhassan, 2022).

2.5 Intrusion Detection Systems

Intrusion Detection Systems (IDS) are crucial for detecting potential cyber threats like Distributed Denial of Service (DDoS) attacks. These systems can be divided into two main categories: signature-based and behavior-based detection systems. Signature-based IDS operate by identifying distinct patterns or "signatures" associated with malicious traffic. They are effective in spotting well-documented threats and can provide detailed alerts when recognized signatures are present (Alhassan, Eze, & Mohammed, 2023). However, these systems have significant drawbacks, especially regarding new or innovative DDoS attacks, as they often fall short in detecting threats that do not have established signatures, commonly termed zero-day attacks (Gorla, Kuo, & Chang, 2021).

Conversely, behavior-based IDS offer enhanced flexibility compared to their signature-based counterparts. These systems analyze network behavior and highlight any activities that diverge from normal patterns (Hajj, 2023). This capability allows behavior-based systems to identify unfamiliar or previously unseen types of attacks. Nonetheless, certain behavior-based IDS might generate false positives, requiring further adjustments to accurately distinguish between legitimate threats and benign activities (Nour & Murtaza, 2020).

The key distinctions between signature-based and behavior-based IDS emphasize the significance of both in contemporary security frameworks. While signature-based systems excel in detecting known threats, behavior-based systems are proficient in monitoring deviations in behavior, making their simultaneous deployment particularly beneficial in defending against DDoS attacks (Khan & Khan, 2021).

2.6 Intrusion Prevention Systems

Intrusion Prevention Systems (IPS) are engineered to block network traffic based on predetermined criteria. These systems often rely on data from Intrusion Detection Systems (IDS) to formulate rules that obstruct malicious traffic (Tao, Ding, & Lin, 2021). Certain IPS solutions can be seamlessly integrated into network traffic flows, enabling real-time detection and blocking of potentially harmful packets, thereby thwarting denial-of-service (DoS) attacks and safeguarding targeted services from excessive traffic volume (Alhassan, Eze, & Mohammed, 2023). Furthermore, IPS act as a vital line of defense against distributed denial-of-service (DDoS) attacks by utilizing flow filters that allow traffic only if it meets specific rules (Hassan, Javed, & Baig, 2022). Both IPS and firewalls can maintain blacklists of particular IP addresses or ranges to mitigate DDoS attack threats (Sajjad & Khan, 2020). Suricata, an open-source IPS, employs signature-based inspection along with real-time filtering capabilities, making it proficient in high-performance threat detection, especially in scenarios experiencing considerable DDoS attacks (Gorla, Kuo, & Chang, 2021). In parallel, cloud-based firewall solutions, such as those offered by Cloudflare and Amazon Web Services (AWS) Shield, bolster security by filtering traffic before it reaches the intended target (Esquivel, Llewellyn, & Khan, 2020). The strategic integration of IPS and firewalls within a holistic framework for DDoS mitigation can significantly reduce the associated risks (Nour & Murtaza, 2020).

2.7 DDoS Mitigation Strategies

Mitigating Distributed Denial of Service (DDoS) attacks can be accomplished through a variety of mechanisms, such as Rate Limiting and Black Hole routing (Kumar, Sharma, & Singh, 2020) (Nour, Murtaza, & Rao, 2022). In addition, preventive measures can be enforced through the implementation of packet filtering rules within firewalls (Alhassan, Eze, & Mohammed, 2023). Over the years, a multitude of strategies have been developed to address the growing complexity and frequency of DDoS attacks (Gupta & Singh, 2020).

2.7.1 DDoS Mitigation Tools

One prominent tool within the cybersecurity landscape is Fail2ban, a security solution engineered to counter brute force attacks, including Denial of Service (DoS) threats. This application functions by analyzing server log files for signs of repeated failed login attempts, allowing it to adjust iptables firewall rules to temporarily block the offending source IP addresses (Alhassan, Eze, & Mohammed, 2023). Although Fail2ban is effective for managing

relatively small-scale attacks, it falls short in providing comprehensive protection against larger Distributed Denial of Service (DDoS) incidents (Gorla, Kuo, & Chang, 2021).

A fundamental feature of Intrusion Detection Systems (IDS) is their capacity to distinguish between normal and abnormal traffic behavior. Tools such as Snort play a critical role in detecting the signatures or traits associated with DDoS attacks. Snort is an open-source tool adept at identifying known attack signatures and notifying administrators during an incident (Nour & Murtaza, 2020). Furthermore, it can be configured to limit specific traffic types, which aids in mitigating organized DDoS attempts. However, signature-based IDS solutions like Snort may demonstrate delays in recognizing new or complex threats, as they depend on established signatures and might miss unique attack patterns (Sajjad & Khan, 2020). As a result, many networks choose to implement Snort alongside behavior-based IDS to improve the detection of unauthorized activities (Khan & Khan, 2021).

2.7.2 DDoS Mitigation Systems

The efficacy of modern cloud-based solutions, such as Akamai and Cloudflare, in countering significant Distributed Denial of Service (DDoS) attacks has become increasingly clear. These services function by utilizing a network of servers that intercept and reroute traffic prior to reaching its intended destination. By capitalizing on this infrastructure, they can effectively mitigate even the most formidable attacks, as demonstrated during the GitHub and Dyn incidents, where traffic was distributed across multiple locations and redirected through their systems. The implementation of cloud-based mitigation services has become a standard strategy for networks seeking to tackle the growing complexity and frequency of DDoS attacks (Nour, Murtaza, & Rao, 2022).

2.7.3 Use of IP Address Reputation in Intrusion Mitigation

Researchers have developed a methodology for profiling malicious IP addresses by analyzing the demographic and socio-economic characteristics of the geographical regions from which attacks are initiated. This analysis involves correlating the locations of attacking clients with the socio-economic conditions at various levels, such as city, state, or national (Husák & Záhonová, 2020). However, a notable limitation of this approach is the assumption that attackers operate from computers located in their residential areas (Esquivel, Llewellyn, & Khan, 2020).

Databases containing known malicious IP addresses are commonly utilized in intrusion prevention efforts. These databases are compiled from reports submitted by organizations that

have experienced attacks and serve various functions, including research reputation scoring and forecasting potential threats (Yang & Alhassan, 2022). The reputation of client IP addresses plays a pivotal role in addressing ongoing threats. In email spam filtering, for instance, a reputation score is assigned to each IP address through a database where email receivers query before approving any requests. Emails sent from IP addresses with low reputation scores are typically rejected, while those from higher-scoring addresses are accepted (Gorla, Kuo, & Chang, 2021).

Moreover, IP address blacklists are employed for web traffic filtering via Intrusion Detection Systems (IDS) like Snort. While these blacklists effectively block IP addresses that continuously engage in malicious behavior, they may inadvertently filter out legitimate addresses and fail to address attacks from new or previously unlisted IPs (Gupta & Singh, 2020). To enhance detection capabilities, some researchers propose a reputation-scoring system that evaluates non-listed IP addresses by considering several factors, including the country of origin, Autonomous System Number (ASN), and Internet Service Provider (ISP) (Tao, Ding, & Lin, 2021). Non-listed IPs that display characteristics like blacklisted ones are flagged as potentially malicious. Blocking these harmful IP addresses can be accomplished using packet-filtering firewalls, while existing DDoS mitigation strategies focus on preventing requests from known malicious sources (Alhassan, Eze, & Mohammed, 2023). Tools, such as Fail2Ban and Iptables, are commonly used to effectively manage the blocking of offending IP addresses (Sajjad & Khan, 2020).

2.8 Global Coordination of IP Addresses and AS Numbers by IANA

The oversight of IP address allocation and Autonomous System Numbers (ASNs) falls under the purview of the Internet Assigned Numbers Authority (IANA) in collaboration with five Regional Internet Registries (RIRs). IANA allocates substantial blocks of IP addresses to these RIRs, which further divide these blocks and distribute them to Internet Service Providers (ISPs) and organizations within their respective regions. The five RIRs consist of the African Network Information Center (AfriNIC), which governs Africa; the American Registry for Internet Numbers (ARIN), responsible for North America; the Asia-Pacific Network Information Centre (APNIC), overseeing the Asia-Pacific region; the Latin America and Caribbean Internet Address Registry (LACNIC), which manages Latin America and the Caribbean; and the Réseaux IP Européens Network Coordination Centre (RIPE NCC), which administers Europe, the Middle East, and Central Asia (Song, Lee, & Lim, 2020). ASNs function as unique identifiers for collections of networks managed by a single administrative entity, such as ISPs

or large organizations, and are crucial for the operation of the Border Gateway Protocol (BGP) routing system, facilitating efficient routing through methods like IP address aggregation and route summarization. RIRs are vital in allocating and managing ASNs on IANA's behalf, ensuring each Internet Autonomous System (AS) is assigned a unique identifier.

WHOIS databases, maintained by various Internet registries and authorities, serve as an essential repository for information related to domain names, IP addresses, ASNs, and their registrants. These databases contain critical data such as the registrant's name, email address, phone number, and physical address, which are publicly accessible to allow interested parties to query and acquire specific information (Gorla, Kuo, & Chang, 2021). WHOIS data is utilized for multiple purposes, including network troubleshooting, legal inquiries, and applications in cybersecurity. It facilitates the identification of domain or IP address owners, the reporting of abuse incidents, the investigation of phishing or spam activities, and the analysis of historical data pertaining to domain names or IP addresses (Khan & Khan, 2021). However, the transparency associated with WHOIS data has raised significant privacy concerns, as the public accessibility of personal information can subject legitimate domain and IP owners to harassment (Nour & Murtaza, 2020). To address this issue, privacy measures, such as obscuring personally identifiable information (PII) from public visibility, have been instituted to protect users while maintaining system functionality (Yang & Alhassan, 2022). For more accurate geolocation data, services like Ipapi provide physical location information without disclosing PII. In this study's context, the key objective is to ascertain the physical location of an IP address at the level of an Autonomous System (AS). Typically, network defense strategies require location data at the country or RIR level, with only a few scenarios necessitating details at the level of an end-user organization. Both WHOIS and Ipapi deliver relevant information, including the AS number, city, and country associated with an IP address; additionally, WHOIS offers insights regarding the range of IP addresses within a network and identifies the corresponding RIR responsible for administering the address (Esquivel, Llewellyn, & Khan, 2020).

2.9 Research Gaps

A range of methodologies has been developed for detecting, mitigating, and preventing Distributed Denial of Service (DDoS) attacks. The deployment of various Intrusion Detection Systems (IDS) provides an effective means for real-time identification of such attacks, which can occur at the host level, across individual networks, or via Content Delivery Network (CDN) providers (Alhassan, Eze, & Mohammed, 2023). Additionally, different configurations of

Intrusion Prevention Systems (IPS) enable the blocking of malicious traffic using firewalls. Tools like Fail2ban and Snort have demonstrated their effectiveness in smaller-scale environments (Sajjad & Khan, 2020). In contrast, CDN providers such as Akamai and Cloudflare present solutions that are suitable for handling larger-scale implementations (Hassan, Javed, & Baig, 2022).

Despite these advancements in DDoS mitigation strategies, notable limitations persist. Tools like Fail2ban and Snort can be sluggish when addressing large-scale attacks, leading to an increased risk of service disruption during major incidents (Gorla, Kuo, & Chang, 2021). Moreover, the necessity of creating a vast number of rules for defensive firewalls ultimately hampers their operational efficiency (Nour, Murtaza, & Rao, 2022). In the event of a suspected DDoS attack, CDN-based defense mechanisms may impose rigorous checks on all incoming requests, which can include the execution of Turing tests for human users. This approach may result in longer processing times and a greater likelihood of legitimate user requests being mistakenly denied (Tao, Ding, & Lin, 2021).

As a result, there is an urgent need for a DDoS mitigation strategy that operates on a granular level by temporarily blocking requests from suspected networks. Such a solution would allow users from unaffected networks to access secured Internet applications without undergoing additional verification processes that could cause delays. This dissertation explores cases in which Internet applications are primarily accessed by users located in regions identifiable by Internet framework structures, such as Autonomous Systems (AS), countries, or Regional Internet Registries (RIR) (Imthiyas & Handan, 2020).

2.10 Conceptual Framework

In the realm of internet services, particularly those involving critical transactions such as e-commerce, online banking, and government platforms, which predominantly serve users situated within specific geographic areas, the proposed approach for mitigating Distributed Denial of Service (DDoS) attacks emphasizes the implementation of strict access controls. This method involves selectively limiting inbound traffic from networks that originate beyond the host country's borders. By focusing on users within the defined service area, who significantly enhance the economic viability and profitability of the digital applications, this strategy seeks to minimize disruptions to legitimate user activities (Kumar, Sharma, & Singh, 2020).

The justification for constraining access to international traffic is based on the premise that limiting foreign requests will have a minimal effect on the overall revenue generated by the services offered. This rationale is supported by the understanding that a substantial percentage

of DDoS attacks originate from sources geographically distant from the targeted systems (Bert, Zhang, & Luo, 2020). Thus, by addressing these external sources, the mitigation strategy not only maintains service availability for local users but also strengthens the financial performance of the applications during periods of increased attack activity.

As illustrated in Figure 2.1, profitability metrics associated with the geographic positioning of users, including factors such as country of origin, Autonomous System (AS) numbers, and Regional Internet Registry (RIR) affiliations—serve as vital indicators. These metrics can inform decision-making processes related to access control measures during DDoS incidents (Li, Liu, & Zhang, 2022). By adopting a data-driven approach that evaluates the reputational status of various network sources, stakeholders can make informed determinations about whether to block or allow traffic from specific locations of interest. This targeted approach aims to enhance the overall resilience of internet applications against DDoS attacks while concurrently protecting the interests of legitimate users and ensuring uninterrupted service continuity. The incorporation of geographic intelligence into DDoS defense mechanisms signifies a proactive step toward developing more robust cybersecurity frameworks that are tailored to address the unique challenges presented by contemporary digital infrastructures (Nour, Murtaza, & Rao, 2022).

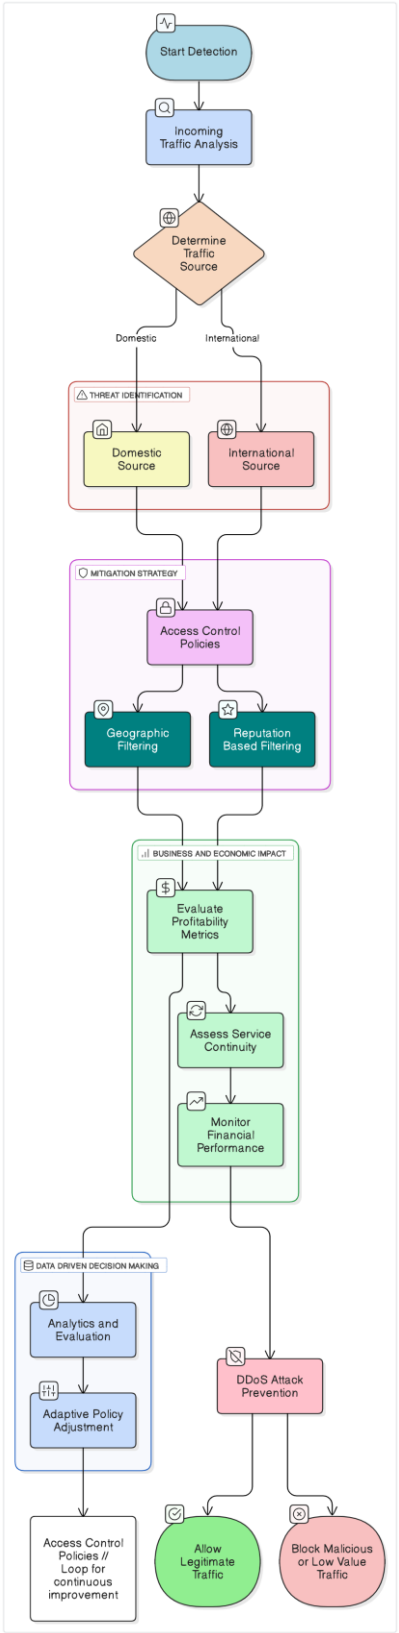

Figure 2.1: Location and Probability Classifier

Chapter Three: Research Methodology

3.1 Introduction

This research examines the efficacy of a novel defense mechanism known as the Multi-Level IP Shedding Defense Mechanism (MLISDM) in mitigating Distributed Denial-of-Service (DDoS) attacks. MLISDM seeks to overcome the limitations present in current mitigation strategies by dynamically reallocating and shedding IP addresses to mislead and deter attackers. This chapter outlines the research design, system development methodology, and ethical considerations involved in the study. It commences with a discussion of how each research question has been addressed.

Research Question 1: What are the predominant patterns observed in Distributed Denial-of-Service (DDoS) attacks, their repercussions, and the strategies employed in recent incidents?

This research question has been explored in the literature review in Chapter 2, which elaborates on the various challenges associated with DDoS attack protection mechanisms.

Research Question 2: How do current DDoS mitigation technologies measure up in terms of efficacy and their inherent limitations?

Addressed in Chapter 2 of the literature review, this question highlights several potential solutions that could enhance the current DDoS mitigation responses.

Research Question 3: What architectural and design factors should be considered when developing a robust multi-level IP shedding defense mechanism for DDoS mitigation?

This inquiry will be elaborated upon in Chapter 4, which will detail the design of the IP shedding DDoS defense scheme.

Research Question 4: In what ways does the multi-level IP shedding defense mechanism differ from conventional DDoS mitigation strategies regarding accuracy, performance implications, and the speed of threat neutralization?

This question will be explored in Chapter 5, where the implementation and testing phases will be discussed in detail.

3.2 Research Design

This study will primarily employ an experimental research design to develop and assess the efficacy of the proposed Multi-Level IP Shedding Defense Mechanism (MLISDM) against Distributed Denial of Service (DDoS) attacks. This approach is chosen for its capacity to create a controlled environment where system variables can be precisely manipulated and their impact on the MLISDM's ability to thwart DDoS incidents can be systematically observed and

measured. The design and objectives of this experimental approach will be substantively informed by an extensive literature review (detailed in Chapter 2), which provides essential conceptual grounding, identifies existing research gaps, and highlights the need for the innovations proposed in this dissertation.

The experimental framework will involve the development of a simulated network environment engineered to accurately reflect key characteristics of real-world internet infrastructure. This setup will enable the replication of various DDoS attack scenarios under regulated conditions. Within this controlled testbed, the MLISDM will be implemented, and its performance and effectiveness will be rigorously evaluated across a range of simulated attack vectors and intensities.

Consistent with an experimental methodology, data collection will be primarily quantitative. The focus will be on capturing objective network performance metrics, including but not limited to the number of client service requests, server response latency, system throughput, the rate of successfully mitigated attacks, and the preservation of service for prioritized user segments.

3.3 System Development Methodology

The design and implementation of the Multi-Level IP Shedding Defense Mechanism (MLISDM) and the associated network simulation environment will utilize an iterative and incremental software development model, particularly emphasizing the Agile Scrum framework. Agile methodologies are well-suited for research projects where requirements can evolve, highlighting the importance of adaptability and responsiveness to changing conditions (Schmidt, Gast, & Jochim, 2021). Scrum provides a methodical approach to dividing development into brief iterative cycles known as Sprints, which promote continuous feedback and refinement throughout the project lifecycle. The development process is structured to encompass several key steps, which will be elaborated upon below:

3.3.1 Requirements Specification

Requirements gathering will be conducted through a review of technical specifications necessary for developing packet filtering rules, criteria for IP pool management, and the underlying logic of the shedding algorithm. Additionally, technical parameters essential for configuring the simulation environment will be assessed.

3.3.2 System Design

After gathering requirements, the system design phase will utilize modeling techniques such as UML diagrams. This process will include Sprint planning to outline specific objectives and deliverables for each iteration. Following this, the coding of the MLISDM modules—encompassing packet filtering, IP management, and the shedding algorithm—along with the features of the simulation environment will be planned accordingly.

3.3.3 System Analysis

System analysis constitutes a vital phase that entails an in-depth evaluation of the existing system's vulnerabilities concerning DDoS attacks, as well as the prospective integration of the recommended defense mechanism. This phase will utilize network analysis tools to delineate the system's topology and pinpoint critical areas that necessitate reinforcement against DDoS threats. Furthermore, the analysis will evaluate the system's performance under normal operational conditions to establish baseline metrics, which will facilitate comparisons following the implementation of the defense mechanism. The examination will also consider the scalability and adaptability of the proposed defense mechanism across diverse system configurations and varied attack scenarios. This aspect is essential to ensure that the defense mechanism can be generalized and applied effectively to a wide range of real-world systems with differing characteristics and requirements.

3.3.4 System Development

The implementation of the DDoS defense system will occur within a testbed composed of Linux computers, configured to operate as both clients and servers. Scripts to simulate attacks will be generated using PHP, Python, and BASH. The enforcement of DDoS defense mechanisms will be created through scripts that employ firewall rules facilitated by tools such as iptables. Version control will be managed through tools like Git.

3.3.5 System Testing

The testing phase will encompass unit testing, wherein the functionality of individual components of the DDoS defense system will be assessed. Additionally, integration testing will be executed to examine the seamless interaction among the combined modules of the DDoS defense system.

Test Environment

The primary testing platform will be a network simulation environment, which will enable the creation of realistic DDoS attack scenarios while maintaining controlled conditions for the

analysis of the Multi-Level IP Shedding Defense Mechanism (MLISDM)'s behavior. Where feasible, real-world attack datasets, such as those from the CAIDA DDoS Attack Datasets (CAIDA, 2023), will be incorporated to enhance the fidelity of the simulations.

Test Cases

A comprehensive suite of test cases will be developed, encompassing a wide array of DDoS attack types, intensities, and durations. This will include Volumetric Attacks, which consist of tests against UDP floods, ICMP floods, DNS amplification, among others. Additionally, protocol-level attacks will assess resilience against SYN floods, Ping of Death, and other exploits at the protocol level. Furthermore, Application Layer Attacks will evaluate the MLISDM's capability to detect and mitigate sophisticated HTTP floods and other application-specific threats.

Performance Metrics

The following performance metrics will be ascertained during system testing: survivability/uptime, attack mitigation rate, resource utilization, and the occurrence of false positives and negatives. Survivability/uptime will measure the system's ability to maintain service availability and essential functions during sustained attacks. The attack mitigation rate will quantify the percentage of malicious traffic effectively blocked or filtered. Throughput will assess the rate at which legitimate traffic is processed during and after an attack. Resource utilization will evaluate the consumption of CPU, memory, and network bandwidth by the MLISDM. Lastly, the false positive/negative metric will gauge the system's accuracy in differentiating between attack traffic and legitimate traffic.

3.4 System Validation

System validation involves the process of confirming that the Multi-Level IP Shedding Defense Mechanism (MLISDM) aligns with its intended specifications, effectively achieves its objectives in mitigating DDoS attacks, and corresponds with the initial research goals. Validation transcends mere testing, aspiring to ensure that the system represents "the right solution" to the identified problem.

Validation Against Requirements: A critical facet of validation will entail mapping the outcomes of system testing back to the functional and non-functional requirements established during the system analysis phase. This evaluation will ascertain whether the MLISDM meets the desired performance levels, scalability, attack mitigation rates, and resource efficiency.

Comparison to Research Goals: The effectiveness of the MLISDM will be assessed against the overarching research questions delineated previously. This assessment will investigate whether the MLISDM demonstrates advancements over existing DDoS mitigation techniques, particularly in addressing the challenges highlighted in the literature review.

Validation Scenarios: In addition to performance-oriented test cases, validation scenarios will be crafted to reflect realistic deployment settings. These scenarios may encompass simulated enterprise network environments, critical infrastructure, or cloud-based services subjected to DDoS attacks. Successful operation of the MLISDM within these contexts will bolster confidence in its practical applicability.

3.5 Ethical Considerations

The design and execution of the Multi-Level IP Shedding Defense Mechanism (MLISDM) necessitates a careful examination of the potential ethical ramifications associated with its operation. Testing and implementation will occur on an isolated network, which will primarily involve physical isolation, though some cases may utilize logical separation. Nonetheless, the following considerations must be addressed for deployment and testing within live networks:

Network Disruption: The process of IP shedding may induce temporary service interruptions for legitimate users, especially during the initial phases of a DDoS attack. The design of the MLISDM will prioritize minimizing such disruptions and facilitate swift restoration of normal services.

Privacy Considerations: The detailed packet analysis and potential IP address logging integral to the MLISDM's operation may raise privacy concerns. Design principles will focus on data minimization, collecting and storing only essential information for attack mitigation, and strict compliance with privacy regulations such as GDPR and the Kenya Data Protection Act of 2019.

Transparency: Operators implementing the MLISDM possess an ethical responsibility to transparently inform users about the possibility of short-term service disruptions and the rationale for employing IP shedding as a defensive mechanism. Open communication fosters trust and ensures users provide informed consent.

Potential for Misuse: While intended as a defensive measure, it is crucial to acknowledge that any robust network manipulation tool carries the potential for misuse. The MLISDM will integrate safeguards to prevent unauthorized access and malicious exploitation.

Fairness and Non-Discrimination: The shedding algorithm of the MLISDM must be meticulously crafted to avert biases that could disproportionately affect specific users or groups based on factors such as network location or traffic patterns.

CHAPTER 4: System Design and Architecture

4.1 System Overview

This chapter provides a comprehensive overview of the design and architectural framework for the Multi-Level IP Shedding Defense Mechanism (MLISDM). The MLISDM is meticulously engineered to dynamically counteract Distributed Denial of Service (DDoS) attacks, with a primary focus on prioritizing business continuity and maximizing revenue for essential online services and critical digital infrastructure. For the purposes of this dissertation and the developed framework, the term 'business,' along with associated concepts such as 'revenue model' and 'critical infrastructure,' is primarily contextualized within, and draws illustrative examples from, business and government operations in Kenya.

Recognizing the increasing sophistication and frequency of DDoS attacks that pose a significant threat to these Kenyan operations, the system is designed to provide robust and adaptive protection. A core principle of the MLISDM is the understanding that not all network traffic holds equal value to the protected entity—envisioned chiefly as a commercial enterprise or a governmental service operating within the Kenyan context. Consequently, the system incorporates a tiered protection strategy directly aligned with the specific entity's revenue model or service criticality. This ensures that its high-value customers and essential operational functions receive preferential treatment and prioritized protection, especially when under duress from an attack. This strategic approach allows for effective DDoS mitigation while minimizing disruption to legitimate, high-value traffic and optimizing operational and financial outcomes for these Kenyan-centric services.

Specifically, the system is designed to handle traffic originating from four distinct regional zones: Zone 1: Service Country (Kenya), Zone 2: Economic Zone (East Africa excluding Kenya), Zone 3: Continent (Africa excluding East Africa), and Zone 4: Rest of the World (excluding Africa). This zonal approach, with Kenya (Zone 1) designated as the highest priority region due to it being the primary service country, facilitates fine-grained control, allows for tailored defense strategies, and enables profitability-aware decision-making for each defined region. The proposed defense mechanism is aimed at protecting online applications whose users are distributed across these four zones, acknowledging that the largest percentage of users typically access such services from Zone 1. In addition to these customer zones, the application is configured with an Access Control List (ACL) of whitelisted critical customer IP addresses. Examples of such protected applications include bank online and API portals, mobile money

payment systems, university eLearning portals, and government services portals, all of which are significant and vital components of Kenya's digital ecosystem and economy.

Zone Description

Zone 1: Service Country (Kenya) - High-value customers: This zone comprises traffic originating from Kenya, representing the highest-value customers. These customers are critical to the business, and their experience is prioritized during an attack.

Zone 2: Economic Zone (East Africa except Kenya) - Significant-value customers: This zone includes traffic from the East African region, excluding Kenya. These customers represent a significant portion of the revenue and would receive a higher shedding priority compared to zones 3 and 4.

Zone 3: Continent (Africa except East Africa) - Medium-value customers: This zone encompasses traffic from the African continent, excluding the East African region. While these customers contribute to the business, their value is lower compared to Zones 1 and 2.

Zone 4: Rest of the World (except Africa) - Low-value customers: This zone includes all traffic originating from outside the African continent. These customers have the lowest value and are subject to the most aggressive IP shedding during an attack.

4.2 Overall System Architecture

The system adopts a multi-layered architecture, illustrated in Figure 4.1, to provide comprehensive protection through a defense-in-depth strategy tailored to different customer segments. This layered approach ensures that multiple lines of defense are engaged, increasing the system's resilience against diverse attack vectors and providing a robust security posture. Each layer plays a crucial role in analyzing traffic, identifying threats, and implementing mitigation strategies, all while considering the business impact of these actions.

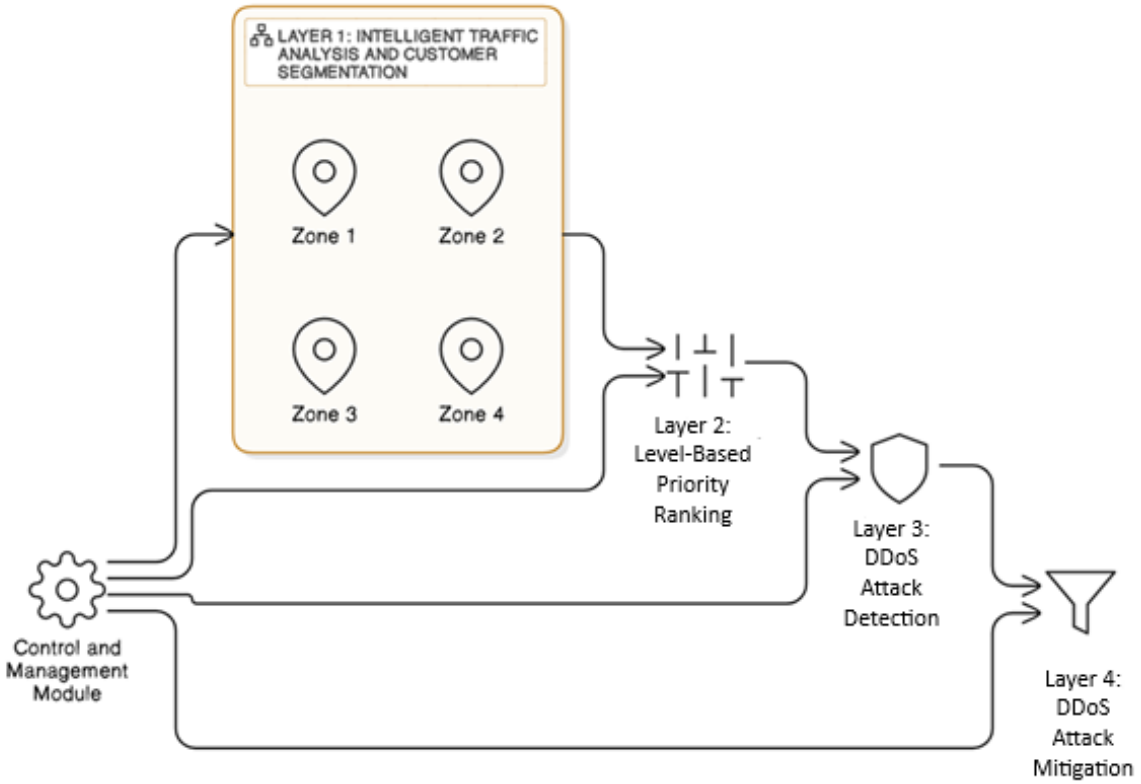

Figure 4.1: System Architecture Diagram

- **Layer 1: Intelligent Traffic Analysis and Customer Segmentation:** This layer offers a set of background processes that run all the time. It serves as the first line of defense, performing real-time analysis of incoming traffic. It leverages both geographic information (obtained through WHOIS lookups) and customer value (determined by analyzing transaction data) to gain a comprehensive understanding of the traffic. This analysis facilitates classification of service request into the four distinct zones; in addition, it updates the zone's revenue contribution by the request.
- **Layer 2: Level-Based Priority Ranking:** Based on the customer segmentation from Layer 1, this layer ranks zones according to number of service requests and revenue contribution over a pre-defined time frame. In the event of an attack, the proposed defense mechanism will aim to maximize the protected application service revenue, by allowing requests from zones with highest contribution to revenue. The mechanism aims to minimize the number of potential users prevented from accessing the application due to IP shedding.

- **Layer 3: DDoS Attack Detection:** This layer is responsible for identifying when the protected system is under a Distributed Denial of Service attack. It is important to note that the design and development of novel DDoS attack detection algorithms themselves are outside the primary scope of this dissertation. The MLISDM architecture is designed to integrate with, and be triggered by, an existing or assumed attack detection capability. This could range from sophisticated behavior-based Intrusion Detection Systems (IDS), machine learning-driven anomaly detection (particularly envisaged for protecting high-value zones), or more straightforward threshold-based alerting mechanisms. Conceptually, Layer 3 continuously monitors network traffic patterns, server load, connection rates, and other relevant system health indicators. When these indicators deviate significantly from established baselines or match known attack signatures, suggesting a DDoS attack is in progress, this layer formally flags an attack state. For the prototype implementation detailed in Chapter 5 (Section 5.6), a pragmatic trigger was utilized to simulate this detection: an attack was considered detected when the number of concurrent incoming connections to the web server exceeded a predefined threshold (e.g., 200 connections), indicating server overload symptomatic of a DDoS event in the controlled testbed. Upon confirming an attack, Layer 3 plays a crucial role in staging the mitigation response. It signals the activation of the subsequent defense layers and provides the necessary context (such as attack severity or type, if discernible by the detection mechanism) to the control and management module. It is at this juncture that the continuously gathered intelligence from Layer 1 (zone identification and customer value) becomes paramount for the targeted response executed by Layer 4. This layer, therefore, acts as the crucial bridge between identifying a threat and initiating an intelligent, value-driven mitigation.
- **Layer 4: DDoS Attack Mitigation:** In the event a DDoS attack has been detected, this layer implements dynamic IP shedding. This involves updating rules in a packet filtering firewall to block requests from IP addresses in low value zones. Shedding thresholds are dynamically adjusted based on the zone revenue value and the severity of the attack, ensuring minimal disruption to highest value zones. As defined earlier, Zone 1 is the highest value zone, however layer 3 processes would identify the highest value zone. This may result in more aggressive shedding for lower-value zones, but it ensures that critical business operations and high-value customers are protected.
- **Control and Management Module:** This module acts as the central brain of the system, providing centralized control and oversight. It allows administrators to define

and modify defense policies, including IP shedding rules, attack detection thresholds, and shedding parameters, adapting the system to evolving threats and business needs. The module also provides real-time monitoring of traffic and attacks, giving administrators a comprehensive view of the network security landscape. Furthermore, it tracks revenue generated by each zone and provides detailed reports on profitability, allowing administrators to monitor the financial impact of defense strategies and make informed decisions.

4.3 IP Shedding Algorithm

IP shedding is a layer 4 defense function. When a DDoS attack is detected by layer 3, the first step in the algorithm is to shed off the lowest value zone. This would be implemented by blocking requests from all IP addresses in the zone, except requests from ACL whitelisted IP addresses. Layer 3 processes will monitor the state of the service over a short Attack Intervention Timeframe (AIT) to determine if the defense action succeeded.

Success is achieved if the arrival rate of requests from permitted zones falls within normal thresholds. In case of successful defense, layer 4 would aim to re-admit the blocked zone(s) at the end of another AIT period. During an AIT period layer 3 IDS monitor the arrival rate of requests for both allowed and blocked zones. Blocked zones would be admitted if the normal state has been attained, otherwise they stay blocked.

In case the first step failed to mitigate the attack, the second lowest value zone would be shed off. State evaluation would happen after another AIT period to determine the outcome of the action. In case of failure, another zone would be shed off and state evaluation performed at the end of AIT period. In case of failure, the last zone would be shed off and state evaluation performed. In case successful defense was achieved at the end of a step in the sequence discussed above, the layer 4 would admit blocked zones in sequence from higher value to lowest value.

ACL whitelisted IP addresses in each zone would be permitted to access the protected application when that zone is shed off. Layer 3 employs rate limiting detection mechanisms to protect the system from attacks originating from compromised clients using ACL whitelisted IP addresses. To meet the goal Component Design of maximizing revenue and minimizing the number of blocked potential users, once a higher value zone gets blocked, a lower value zone may be admitted as long as it meets rate limiting thresholds measured by layer 3. The operation of the IP shedding algorithm is illustrated in Figure 4.2.

4.4 Data Flow and Interactions

The overall operation of the Multi-Level IP Shedding Defense Mechanism (MLISDM), illustrating the interplay between its core components, is depicted in Figure 4.2 (Sequence Diagram - IP Shedding Process) and Figure 4.3 (Data Flow Diagram). The following narrative further clarifies the data flow, particularly the relationship between zone identification, customer value assessment, and the attack detection component leading to mitigation.

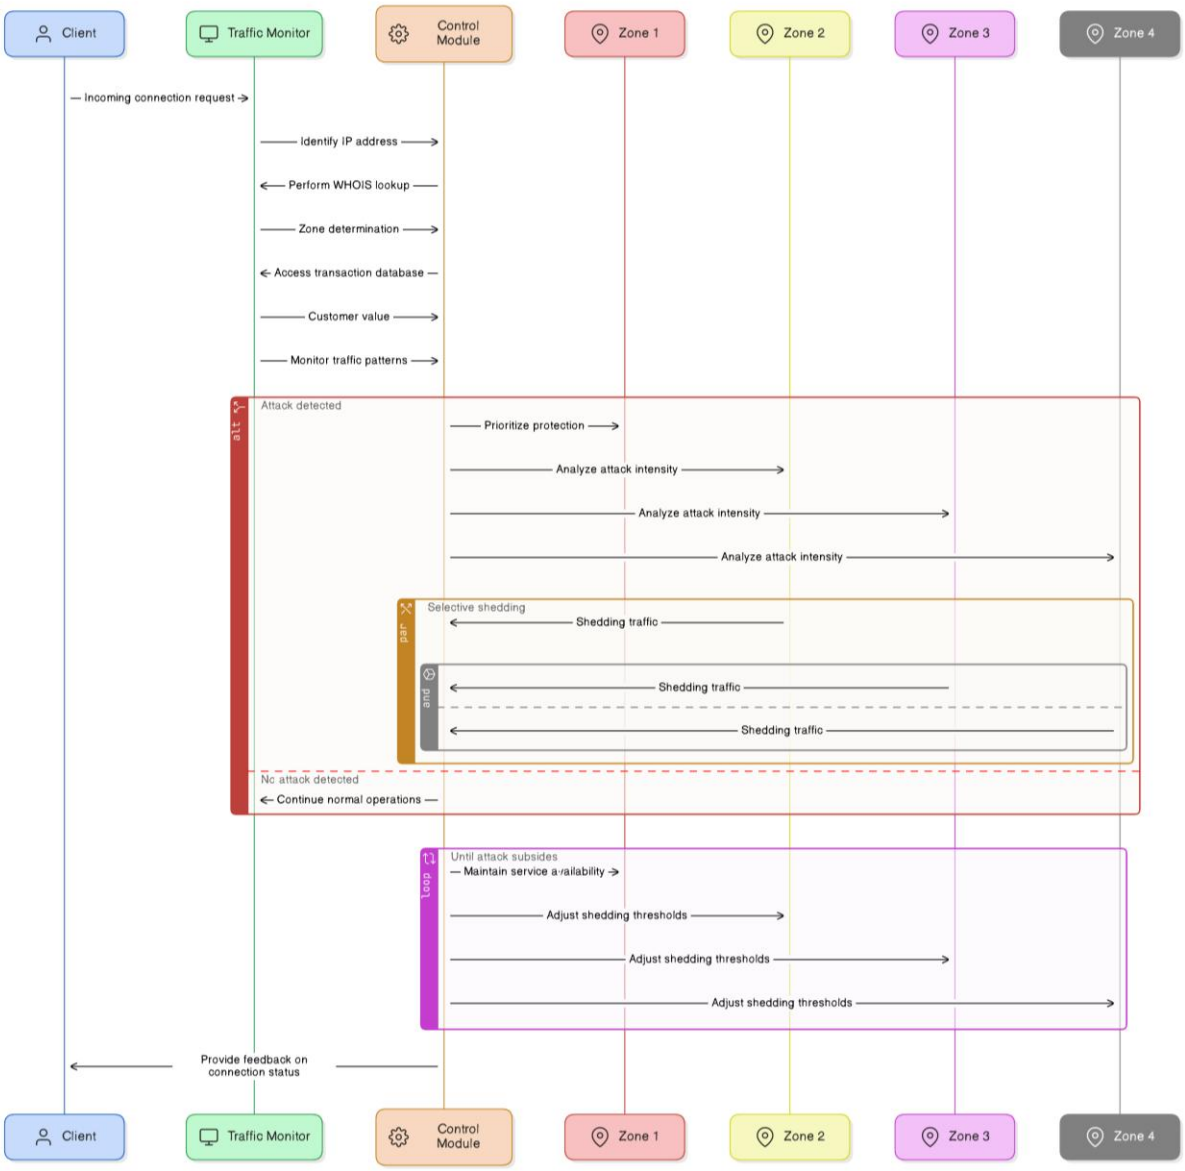

Figure 4.2: Sequence Diagram – IP Shedding Process

1. Continuous Traffic Classification (Layer 1 Processing):

   As incoming connection requests arrive at the system, they are immediately subjected to processing by Layer 1: Intelligent Traffic Analysis and Customer Segmentation.

- The IP Address Location Classifier component performs zone determination. It identifies the source IP address and, through WHOIS lookups or geolocation services (as implemented with ip-api.com in the prototype, Section 5.4), maps it to its appropriate predefined geographic zone (e.g., Zone 1: Service Country - Kenya).
- Simultaneously, the Profit Classifier component (Section 5.5) assesses the customer value associated with the traffic. This is typically based on the source zone's aggregated historical transaction data or predefined value tiers. This classification process occurs continuously, providing the MLISDM with an ongoing, real-time understanding of traffic origins and their associated business importance. This rich contextual data is stored and readily available.

2. Attack Detection and Confirmation (Layer 3 Functionality):

   In parallel with Layer 1's continuous classification, the Attack Detection component (conceptually residing within Layer 3, or represented by the server load trigger in the prototype as per Section 5.6) monitors the overall system state. This involves observing aggregate traffic patterns, connection limits, server resource utilization, or other indicators for anomalies that signify a potential DDoS attack.

   The "Attack Detection" step, as illustrated in Figure 4.3, can be understood as an initial system-wide check or a trigger point. If this initial monitoring indicates a potential attack (e.g., server connection threshold exceeded), it then more formally engages the full analytical capabilities of Layer 3 to confirm the attack state.

3. Integrated and Intelligent Mitigation Response (Layers 2, 3, and 4):

   Once Layer 3 confirms that a DDoS attack is actively underway:

   - The system immediately leverages the comprehensive contextual information (Zone ID and Customer Value) that has been continuously established and updated by Layer 1 for all relevant traffic flows.
   - Layer 2: Level-Based Priority Ranking utilizes this zone and value information to dynamically rank the defined zones based on factors like their current revenue contribution and the volume of legitimate service requests. This ranking informs the shedding priority.
   - Layer 3, having confirmed the attack, signals the Control and Management Module and Layer 4: Dynamic Packet Shedding to initiate mitigation.

- Layer 4 then executes the IP Shedding Algorithm (detailed in Section 4.3). This algorithm's decisions are highly informed and precise because it directly references the pre-established Zone Identification and Customer Value attributes of the various traffic segments (or entire zones). It will then strategically proceed to shed traffic, typically starting with zones identified as having the lowest business value and/or posing the highest current attack risk.

The data flow is thus interactive and adaptive: Layer 1 provides continuous, granular classification of all traffic. Layer 3 monitors the system for attack indicators and confirms an attack state. Upon confirmation, the intelligence from Layer 1, prioritized by Layer 2, is used by Layer 4 to execute a precise, value-driven IP shedding response, minimizing disruption to the most critical services and user segments. Figure 4.3 illustrates this overall flow, from initial IP identification and classification through to the selective shedding based on combined intelligence when an attack is detected.

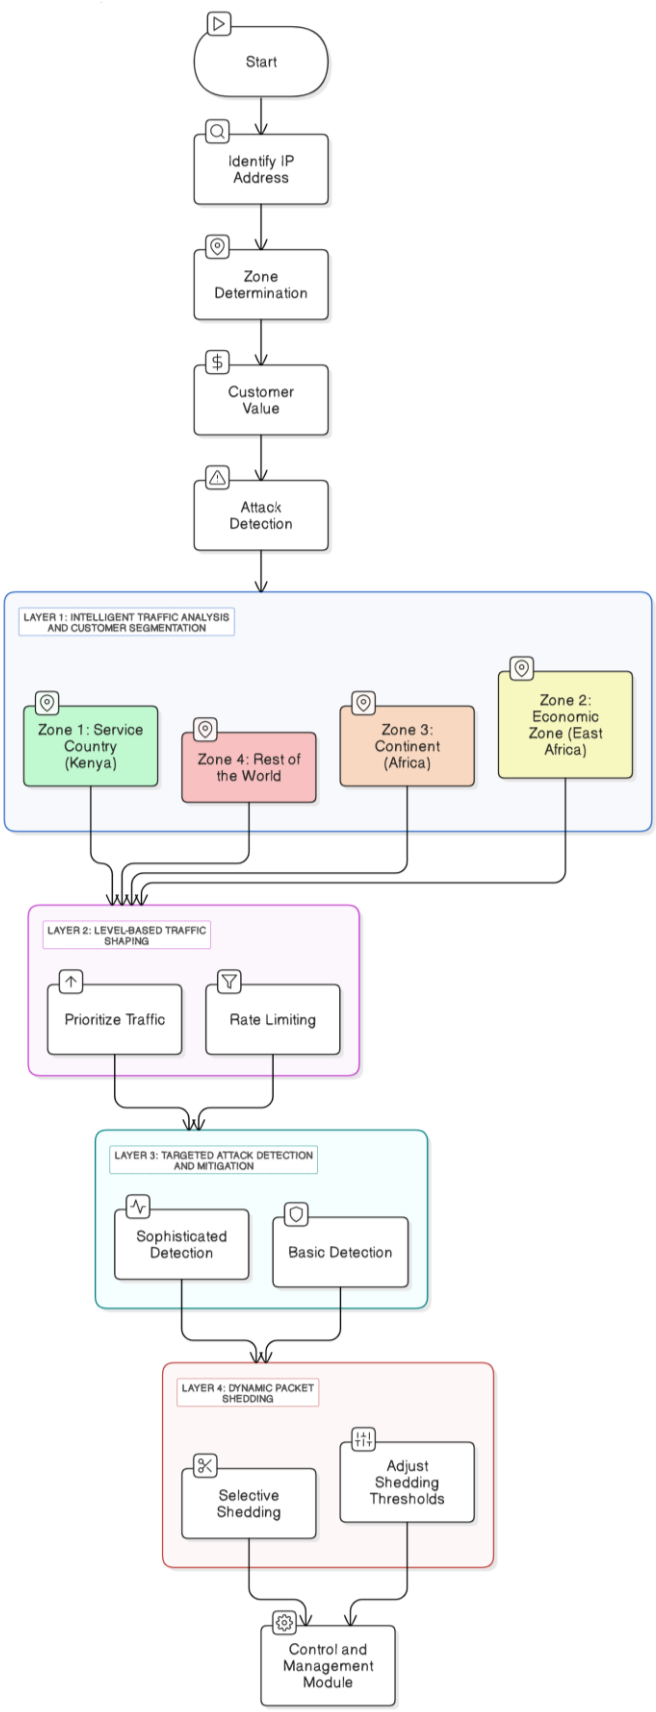

Figure 4.3: Data Flow Diagram

Chapter 5: Implementation And Testing

5.1 Introduction

This chapter presents the implementation and testing of the proposed IP Shedding DDoS defense mechanism. It starts with implementation and testing of individual components of the integrated system, followed by the integrated system. The implementation is an emulation of the real Internet environment. For this reason, it replicates the behavior of real Internet systems albeit in a controlled environment.

5.2 System Environment

Components of the defense system were implemented on the Debian Linux version 11 operating system. Debian is an open-source operating system that offers flexibility to developers with respect to achieving custom configurations that were necessary for this implementation. It allows the installation of various user-space applications. The following application software was used in this implementation: apache2 web server (Apache Software Foundation), iptables firewall (Netfilter Project), and MariaDB database (MariaDB Foundation). In addition, the following scripting and programming languages were used: PHP server-side web programming language, and BASH shell scripting language.

The setup comprised a Debian virtual machine hosting one web server acting as the protected application server, and two client virtual machines. The Virtual Machines (VM) were deployed on the Oracle virtual box hypervisor (Oracle).

5.3 Testbed Setup

A. Operating System

Debian Linux release 11 (bullseye) was installed on Oracle VirtualBox VM installed on Microsoft Windows 11. The system specifications were defined on virtual box as follows: 64-bit OS, 1024 MB RAM, 20 GB hard disk storage space. The network was set to bridged adaptor and mapped to the IEEE 802.11 (WLAN) interface of the host computer. Bridged adaptor setting enabled each VM to communicate on the test network using its IP address, without having to use an NAPT gateway. The two additional VMs were cloned from the primary VM.

B. Web server

The Apache2 web server, version 2.4.62, was provisioned on the Debian VM1 using the Advanced Packaging Tool (APT). This method ensured efficient dependency

resolution and simplified maintenance. The installation command, executed with superuser privileges, is:

```
sudo apt update && sudo apt install apache2=2.4.62-1
```

Where: sudo for elevated privileges, apt update to synchronize the local package index, && to ensure sequential execution, and apt install apache2 to initiate the installation process. APT handles the downloading, unpacking, installation, and basic configuration of the Apache2 package and its dependencies.

To support the execution of PHP applications, both Apache-specific and command-line interface (CLI) PHP modules were installed. The Apache module integrates PHP with the web server, allowing it to process PHP code within web pages. The CLI modules enable PHP script execution from the terminal, which was particularly useful for testing and development purposes. The installation command is:

```
sudo apt install php libapache2-mod-php php-cli
```

If needed, additional PHP modules can be installed as needed using “sudo apt install php-<module-name>”. After installation, Apache was reloaded for changes to take effect:

```
sudo systemctl restart apache2
```

The web server was configured by setting the maximum number of tasks (processes or threads) that the Apache service is allowed to create to 200. The value can be displayed using the command: systemctl show apache2 | grep TasksMax. The maximum number of child processes is enforced by systemd. This translates to the maximum concurrent client requests the web server can accept. In the event of a DDoS attack, the number of incoming requests is bound to be multiple times higher than this value. This will overload the server and prevent it from accepting requests from legitimate clients.

The TasksMax setting (which directly relates to settings like MaxRequestWorkers or MaxConnections within Apache's configuration, depending on the MPM) acts as a crucial defense mechanism against DDoS attacks. When the number of incoming requests exceeds 200, any further requests will be queued or rejected, depending on the specific configuration and the nature of the underlying operating system's TCP stack.

This is a critical consideration: although the server will continue to receive a surge of requests, it will refrain from processing all of them concurrently. This strategy safeguards the server from depleting its resources (such as CPU, memory, and network

bandwidth) and address an excessive number of malicious requests. As a result, legitimate clients may experience some delay, but they have an increased likelihood of having their requests successfully processed since the server is not entirely overwhelmed.

It is crucial to recognize that this approach does not serve as a complete solution to Distributed Denial-of-Service (DDoS) attacks. Rather, it functions as a mitigation strategy. While it helps prevent the server from becoming unresponsive, it does not eliminate the attack entirely. A well-executed DDoS attack can still overwhelm the network bandwidth, thereby obstructing legitimate traffic from reaching the server, even if the Apache software remains operational. An effective DDoS defense framework necessitates a multi-layered strategy.

C. Database Server

The database server was created for storing zone data during the location classification tests. Tests in this dissertation were not aimed at the optimal setup of the database, hence details such as the Entity Relationship Diagram (ERD) are omitted. The database server was provisioned by installing MariaDB server version 10.5.28. The default limit on child processes was 7257. This value was left unchanged since the DDoS tests were confined to overloading the web server.

D. Network Setup

The tests were conducted on a private network, thus private IP addresses were assigned to the VMs. The network used class B private IP addresses (Rekhter, Moskowitz, Karrenberg, de Groot, & Lear, 1996) e.g. 172.16.13.x, where x was unique to each VM. The web server's IP address was 172.16.13.10. The clients were assigned IP addresses 172.16.13.20 and 172.16.13.30 respectively.

5.4 IP Address Location Classifier Component

The IP Address Location Classifier component is a critical element of Layer 1 (Intelligent Traffic Analysis and Customer Segmentation) within the Multi-Level IP Shedding Defense Mechanism (MLISDM). Its primary function is to determine the physical geographic origin of incoming IP traffic, which is then used to assign the traffic to predefined operational zones. This zoning is fundamental to the MLISDM's ability to apply targeted, value-driven defense strategies.

The classifier utilizes a public IP geolocation database, accessed via an Application Programming Interface (API) request to a service such as https://ip-api.com (as implemented in the prototype). For any given public IPv4 address, this service can return various data fields, including the country, city, region, latitude, longitude, Internet Service Provider (ISP), Autonomous System Number (ASN), and the associated Regional Internet Registry (RIR). For the MLISDM's zoning logic, the physical location (country, city) is the most critical piece of information extracted.

It is important to understand the distinction between an IP address's physical geolocation and its RIR. The RIR (e.g., ARIN, RIPE NCC, AfriNIC, APNIC, LACNIC) indicates the organization that was originally allocated that specific block of IP addresses. Due to the global nature of internet operations, legacy allocations, or assignments by multinational entities, an IP address block originally allocated by one RIR (e.g., ARIN for North America) might be used to provide services in a geographical region typically covered by another RIR (e.g., Europe served by RIPE NCC, or Africa served by AfriNIC). Therefore, the RIR listed for a specific IP address reflects its registration and allocation history, which may not always directly correspond to the RIR for the geographical region where the IP address is currently in use. The MLISDM's zoning mechanism, however, prioritizes and relies on the determined current physical location of the IP address for its operational decisions regarding zone assignment.

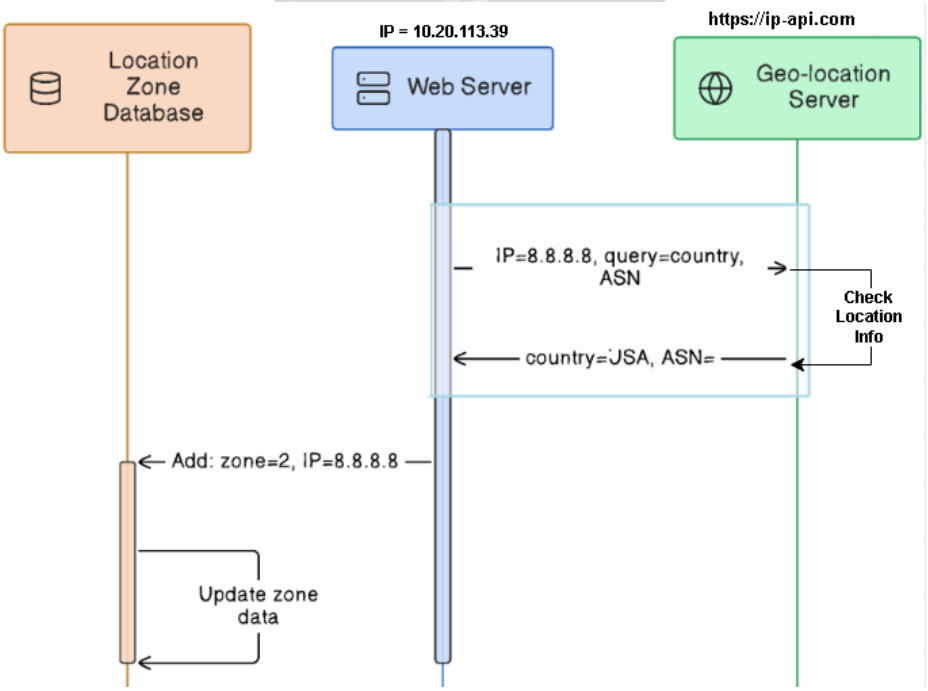

Figure 5.1: IP Location Classifier Component

For the experimental setup and testing detailed later in this dissertation (Section 5.7), a simplified two-zone model was adopted for the emulation platform to demonstrate the core IP shedding functionality. In this model:

- Zone 1 was configured to represent IP addresses originating from within Kenya, designated as the primary service country and the main target for protection.
- Zone 2 was configured to represent IP addresses originating from outside Kenya (i.e., international locations). For simulating attacks, the attacker's traffic in the testbed was set to originate from this Zone 2.

Figure 5.2 illustrates an example of the data fields that can be retrieved and stored by the system in a MySQL database table for IP zone information. The 'Zone' assignment in this table (e.g., Zone 1 or Zone 2) is determined based on the physical 'Country' field, aligning with the experimental setup described above. For instance, an IP geolocated to Kenya would be assigned to Zone 1, while an IP geolocated to a country outside Kenya (like Ireland in the example) would be assigned to Zone 2. The RIR information, while recorded, is secondary to the physical location for this zoning process.

	id	ipAddress	zone	country	city	asn	ISP	LIR	RIR
e	1	34.243.183.166	2	Ireland	Dublin	AS16509 Amazon.com, Inc.	Amazon Technologies Inc.	AWS EC2 (eu-west-1)	ARIN (North America)
e	2	156.0.233.51	1	Kenya	Nairobi Hill	AS328225 Strathmore University	Strathmore University	Strathmore University	ARIN (North America)

Figure 5.2: IP Zone Details Table

Note: The RIR (Regional Internet Registry) listed reflects the original allocation body for the IP address block and may not correspond to the RIR for the geographical region of the IP's current use. For example, an IP address geolocated to Kenya (AfriNIC region) or Ireland (RIPE NCC region) might show ARIN as its RIR if the block was originally allocated by ARIN to a global entity. The MLISDM's zone assignment is based on the physical 'Country' of the IP address.

The general context for IP address administration in Africa is that the African Network Information Centre (AfriNIC) is the RIR responsible for allocating IP address resources. Organizations within Kenya, such as Internet Service Providers or institutions like Strathmore University, typically obtain their IP address blocks from AfriNIC, sometimes via Local Internet Registries (LIRs) like the Kenya Network Information Centre (KeNIC), which also manages the .ke ccTLD. However, as explained, specific IP addresses in use within Kenya might still

trace their ultimate RIR registration back to another body like ARIN due to the history of their allocation.

A sample of the PHP source code used for the IP location determination and data storage is provided in Appendix C. This script queries the ip-api.com service and populates the local database with the relevant geographical and network information, which is then used by the MLISDM for its decision-making processes.

5.5 Profit Classifier Component

The Profit Classifier Component, integral to Layer 1 (Intelligent Traffic Analysis and Customer Segmentation), is designed to assess the financial or business value associated with incoming traffic. This valuation is crucial for the MLISDM to make informed decisions about traffic prioritization and selective IP shedding during a DDoS attack, aligning defensive actions with business continuity and revenue preservation objectives.

5.5.1 Conceptual Data Sources and Integration for Profit Classification

In a real-world deployment, the Profit Classifier component would ideally integrate with various backend systems to obtain transaction and revenue data. The specific sources would depend on the nature of the protected application:

- E-commerce Platforms/Online Retail: Integration with order management systems or payment gateways could provide data on transaction values, customer purchase history, and average order value. This data could be linked to user accounts, which in turn might be associated with IP addresses or, more broadly, with the geographic zones from which users typically access the service.
- Financial Services (e.g., Online Banking, Payment Portals): Data could be sourced from transaction processing systems, indicating the volume or value of financial transactions. For business clients, this might involve API usage metrics that correlate with service fees.
- Subscription Services/SaaS Platforms: Customer Relationship Management (CRM) systems or billing platforms could provide information on subscription tiers, customer lifetime value, or service usage levels that correspond to different revenue contributions.
- Government Service Portals: While direct "profit" may not apply, "value" could be defined by the criticality of the service, the user base served (e.g., essential services for

citizens vs. informational queries), or metrics indicating the successful completion of vital civic processes.

- Application Logs: For many applications, specific user actions that signify high value (e.g., completing a checkout, accessing a premium feature, performing a critical API call) can be logged with associated session/user identifiers. These identifiers can then be correlated with IP addresses or user zones.

The core idea is to establish a quantifiable measure of "value" or "profitability" associated with different segments of traffic. This data would then be aggregated and attributed to the respective geographic zones from which the traffic originates. For instance, the system could maintain a dynamic profile of the average transaction value or revenue per user originating from Zone 1 versus Zone 2 over a given period.

5.5.2 Data Scanned/Used

The data used by the Profit Classifier would typically include:

Transaction Identifiers and Values: Monetary value of purchases, service fees, subscription payments.

User/Customer Identifiers: To link transactions to specific users or accounts.

Session Information: Including source IP addresses, which are then mapped to zones by the IP Address Location Classifier.

Timestamps: To analyze revenue patterns over time and maintain up-to-date profitability scores for zones.

Service/API Usage Metrics: For B2B services, the volume or type of API calls can indicate value.

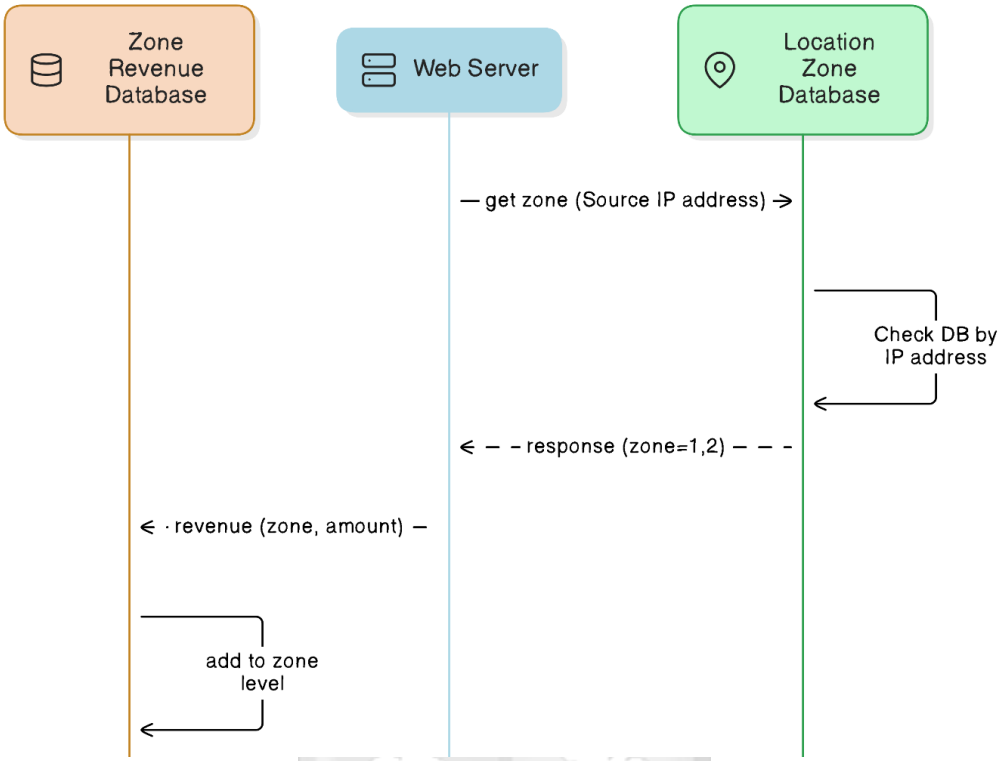

Figure 5.3: Profit Classifier component

Figure 5.3 illustrates this conceptual data flow. When a web server receives a request that results in a value-generating event (e.g., a purchase, a critical transaction):

1. The web server (or an integrated application-level monitor) identifies the source IP address of the request.
2. It queries the "Location Zone Database" (maintained by the IP Address Location Classifier) to determine the zone of the source IP.
3. The value of the transaction/event (e.g., "amount") is captured.
4. This value, along with the determined zone, is then sent to the "Zone Revenue Database" (or a similar data store managed by the Profit Classifier).
5. The Profit Classifier component updates its internal metrics, such as the cumulative revenue or average transaction value for that specific zone. This updated "zone profit level" is then available to the MLISDM for decision-making.

5.5.3 Prototype Implementation (Simplified Emulation):

For the prototype implementation detailed in this dissertation, a simplified mechanism was employed to emulate the core logic of the Profit Classifier. As shown in the client request tool (Appendix F) and the web server application code (Appendix G):

1. Client requests were programmed to send a randomized value parameter (e.g., ?value=100) with each request to the web server.
2. The server-side PHP script (app.php) received this value.
3. The script then determined the client's zone based on its IP address (using a simplified hardcoded mapping for the testbed IPs: 172.16.13.20 to Zone 1, 172.16.13.30 to Zone 2, as per Section 5.7 and Appendix G's getZoneVal function).
4. This value was then logged against the determined zone in the revenueData table and used to update the aggregate value for that zone in the profitabilityData table (Appendix D shows the conceptual PHP for profit classification, and Appendix G shows the server application logic).

This simplified approach allowed the prototype to simulate the dynamic association of "revenue" with different zones, enabling the testing of the MLISDM's value-driven IP shedding logic without requiring complex integrations with real backend financial systems. It is acknowledged that this is an emulation, and a production system would require more robust data sourcing as described conceptually above.

5.6 DDoS Integrated Defense Mechanism

The DDoS Integrated Defense Mechanism represents the core operational logic of the MLISDM, bringing together the intelligence gathered by its foundational components to execute an adaptive and value-driven response to detected attacks. This mechanism relies on the continuous, real-time data provided by the IP Address Location Classifier component (Section 5.4), which identifies the geographic origin of traffic and assigns it to predefined zones, and the Profit Classifier component (Section 5.5), which assesses the business value or revenue contribution associated with these zones.

The integration of these components is crucial: the IP Location Classifier provides the "where," and the Profit Classifier provides the "how valuable." When an attack is detected, this combined intelligence allows the MLISDM to move beyond generic blocking and implement a nuanced defense that prioritizes critical services and high-value user segments.

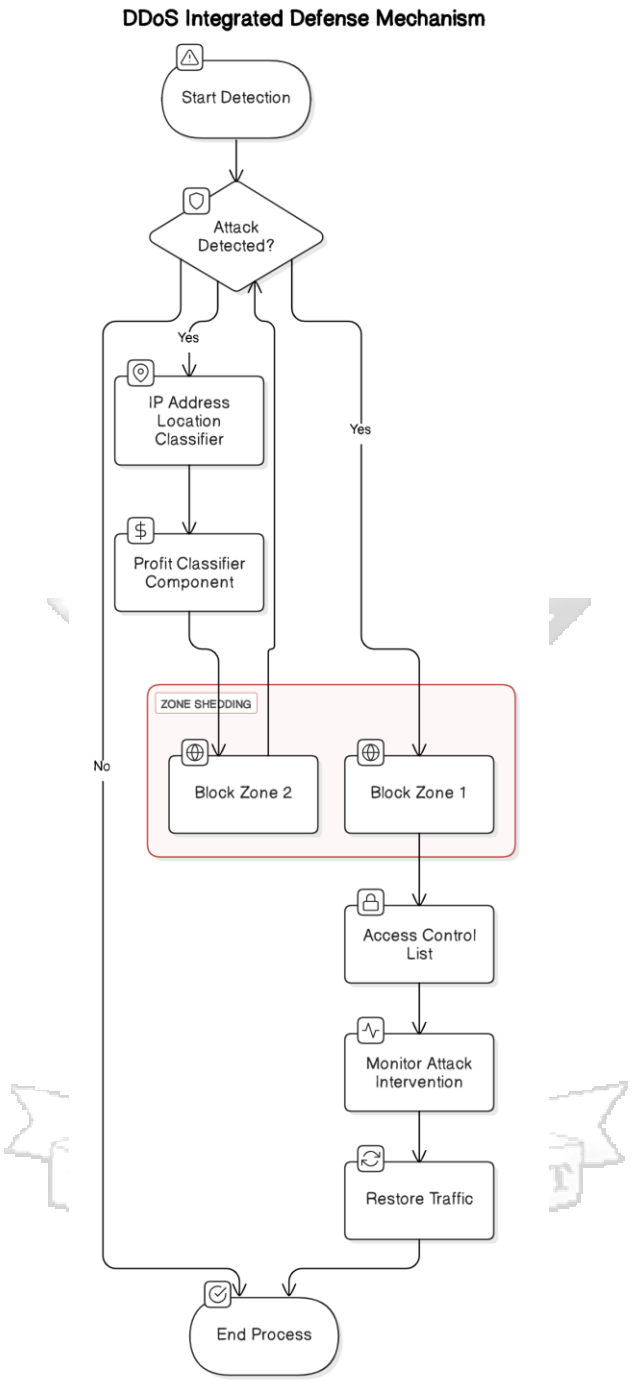

Figure 5.4: DDoS Integrated Defense Mechanism

Figure 5.4 illustrates the decision-making process of this integrated system. The logical flow, based on the clarified inputs from the preceding components, is as follows:

Attack Detection Trigger: The process initiates when an attack is detected. As discussed in Section 4.2 (Layer 3) and for the prototype in this chapter (Section 5.3.B and later in 5.7), this detection can be based on various indicators, such as exceeding server connection thresholds

(the trigger used in the prototype, set at 200 concurrent connections), or more sophisticated IDS alerts in a production environment.

Leveraging Classified Intelligence: Upon attack detection, the system immediately utilizes the data from:

- IP Address Location Classifier: To understand the zonal distribution of current traffic, including potentially malicious flows.
- Profit Classifier Component: To understand the current revenue/value ranking of these zones.

Zone Shedding Logic: The system then implements selective zone shedding based on a pre-defined strategy that aims to maximize service availability for high-value zones while mitigating the attack. In the simplified two-zone testbed:

- If an attack is detected and Zone 2 (International/lower-value/attacker's source in tests) is implicated or is the designated first-to-shed zone, traffic from Zone 2 is blocked via firewall rules (Appendix H). This is the "Block Zone 2" path in Figure 5.4.
- If the attack persists or escalates, or if Zone 1 (Kenya/higher-value) is also significantly compromised or targeted by distinct attack vectors, the system might then also restrict traffic to Zone 1 ("Block Zone 1" path), potentially allowing only whitelisted critical IPs through (as per the "Access Control List" step). This represents a more severe defense posture.

Continuous Monitoring and Restoration: Throughout the mitigation, the system monitors the attack status ("Monitor Attack Intervention"). If the attack subsides, the system can gradually restore traffic from blocked zones ("Restore Traffic"), typically starting with higher-value zones.

This integrated approach, underpinned by clear data from the location and profit classifiers, ensures that the defense mechanism is not only reactive but also intelligent, aligning its actions with the strategic goal of protecting the most valuable aspects of the online service. The clarity of these foundational components (5.4 and 5.5) is therefore essential for the robust functioning and logical integrity of the integrated defense mechanism illustrated in Figure 5.4 and tested subsequently

5.7 Testing of Defense Mechanism

Testing was done using two zones, with one client VM per zone. Zone 1 was set up to be the highest revenue zone. The attacker was set up to launch from zone 2; it was hosted on the same VM as the legitimate client of the zone. The web server was hosted on a separate VM. Baseline tests were conducted to determine the performance of the protected web server, as well as clients before a DDoS attack. Subsequently, the performance of the server and clients during a DDoS attack was recorded. During an attack, the Slowloris tool was used to send a high specified number of requests to overwhelm the web. Connections that exceed the maximum number 200, set in section 5.3 will overload the server and prevent some legitimate requests from being accepted.

Each set of tests was done over a period of 5 minutes i.e. 300 seconds

5.7.1 Server Load Measurement Tool

Two server load parameters were used to measure the load. The first parameter was the number of established HTTP sockets representing the volume of requests received by the web server. The second parameter was the number of active Apache HTTP worker processes. Appendix E shows the BASH program for measuring the server load. Samples of the server load were taken at intervals of 1 second. This means the accuracy of measured server load was not 100%.

5.7.2 Slowloris DDoS Attack Tool

Slowloris (Akamai, 2022) is a sophisticated DDoS attack tool that employs a "low and slow" approach to overwhelm web servers. Unlike traditional DDoS attacks that rely on high volumes of traffic, Slowloris operates by opening multiple connections to a target server and sending partial HTTP requests without completing them. This technique keeps connections open indefinitely, gradually exhausting the server's connection pool until it can no longer accept legitimate requests.

The tool was created by Robert "RSnake" Hansen and is named after the slow-moving Asian primate, reflecting its methodical approach to disabling web servers. Slowloris has proven particularly effective against Apache 1.x and 2.x web servers, which have limited connection pools designed for quick request completion.

What makes Slowloris especially dangerous is its ability to operate with minimal bandwidth requirements. A single computer running Slowloris can potentially take down a high-profile server, making it a popular tool for hacktivism. Additionally, Slowloris can be configured to

send different host headers and suppress log file creation during attacks, allowing it to evade detection by conventional security monitoring systems.

The attack works by exploiting a fundamental aspect of the HTTP protocol: the requirement for requests to be terminated by a sequence of newline characters. By withholding these termination sequences, Slowloris forces servers to keep connections open and resources allocated while waiting for completion. As these incomplete connections accumulate, the server eventually reaches its connection limit and becomes unresponsive to legitimate users.

The command for launching an attack with Slowloris on Debian 11 is:

python3 slowloris.py -s 500 -p 80 -v <target-server-ip>

Where:

-s 200 → Number of concurrent connections (adjust as needed)

-p 80 → Target port (default Apache port)

-v → Verbose output

Slowloris Baseline Test

The web server was attacked with 500 concurrent connections. This exceeded the maximum 200 child processes that were configured for the web server. The results in Fig. 5.5 show the number of connections received by the web server, as well as the number of active successful connections accepted by the server. Requests sent by Slowloris were not processed by a dedicated web application.

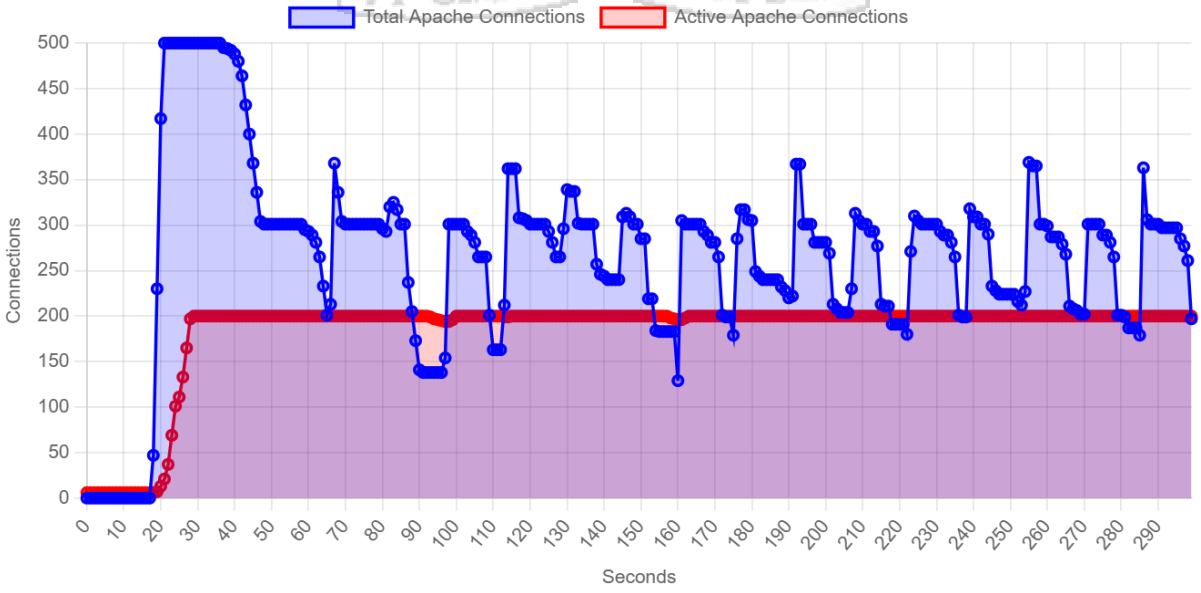

Figure 5.5: Web server performance for Slowloris baseline test

5.7.3 Client Request Tool

Client requests were generated using a custom PHP script that sent either 100 or 200 requests in 5 second intervals. The client recorded the number of successful as well as failed requests. The source code is presented in Appendix F. A PHP application was installed on the web server to receive and process client applications. It recorded the IP zone and transaction value for each client request in a MySQL database. The source code of the web server application is given in appendix G.

5.7.4 Client Baseline Tests

Tests were conducted to determine the server load, and the client performance when one or more clients sent requests to the web server. Client performance comprised the number of successful requests, and the number of failed requests. A DDoS attack was not launched during these tests.

One Client

Figure 5.6 shows server load with one client sending 100 requests at 5 second intervals. While Figure 5.7 shows the number of successful and failed client requests.

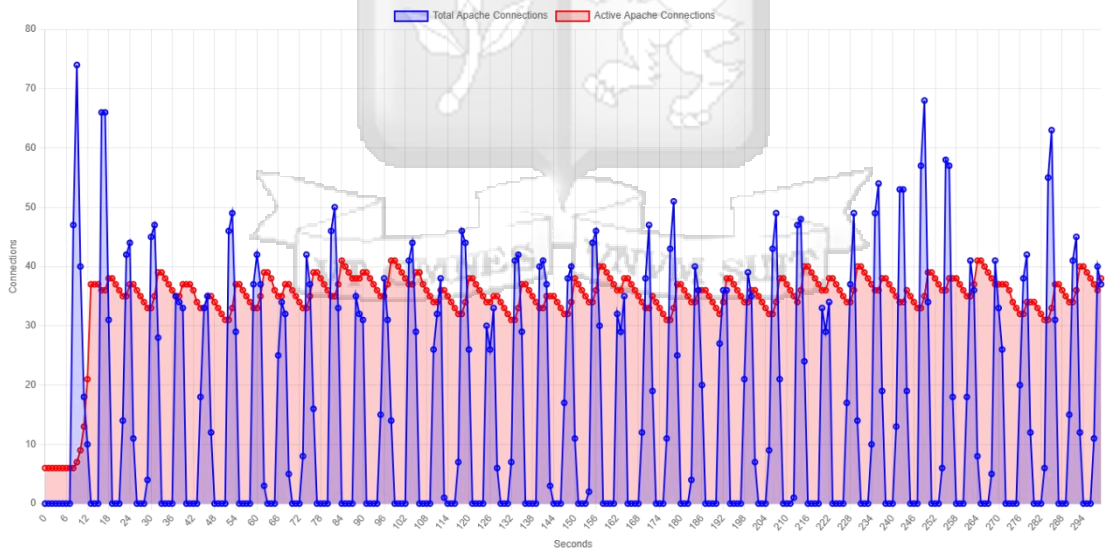

Figure 5.6: Server load with one client – 100 requests every 5 seconds

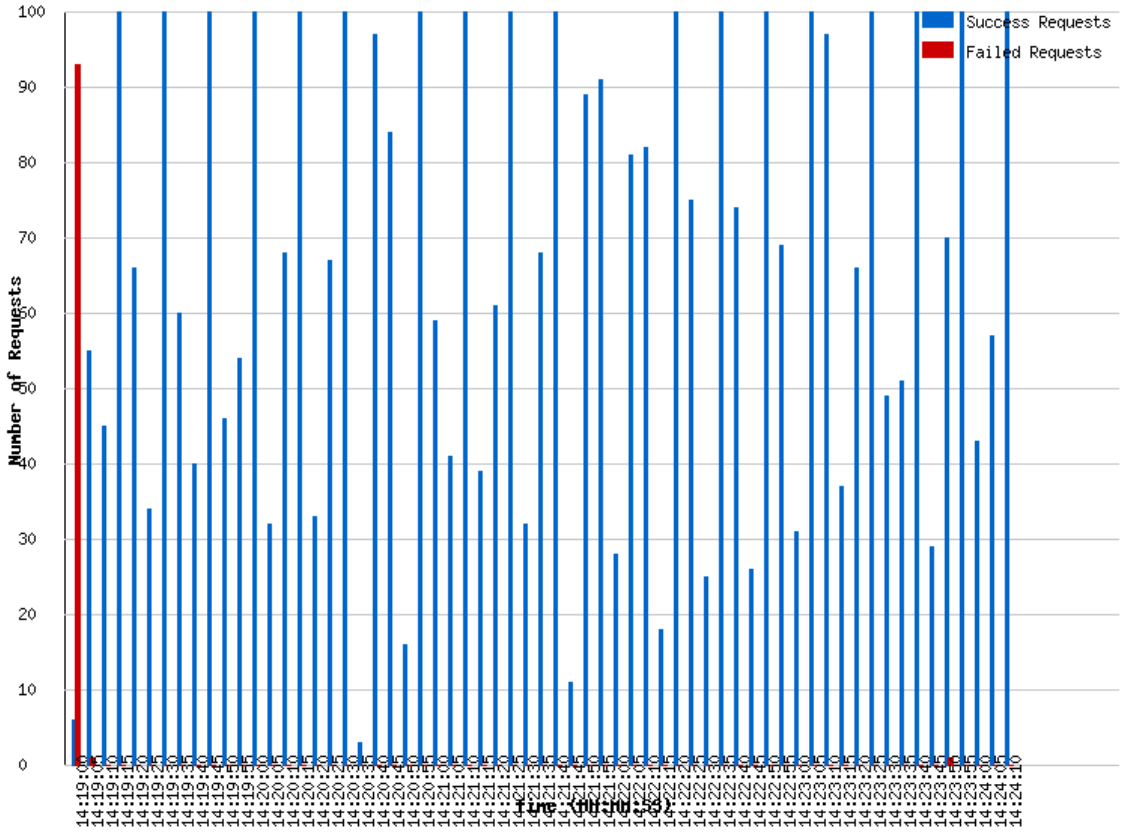

Figure 5.7: Client Performance – One client at 100 request every 5 seconds

Two Clients

Figure 5.8 shows the server load when two clients were each configured to send 100 requests in 5 second intervals.

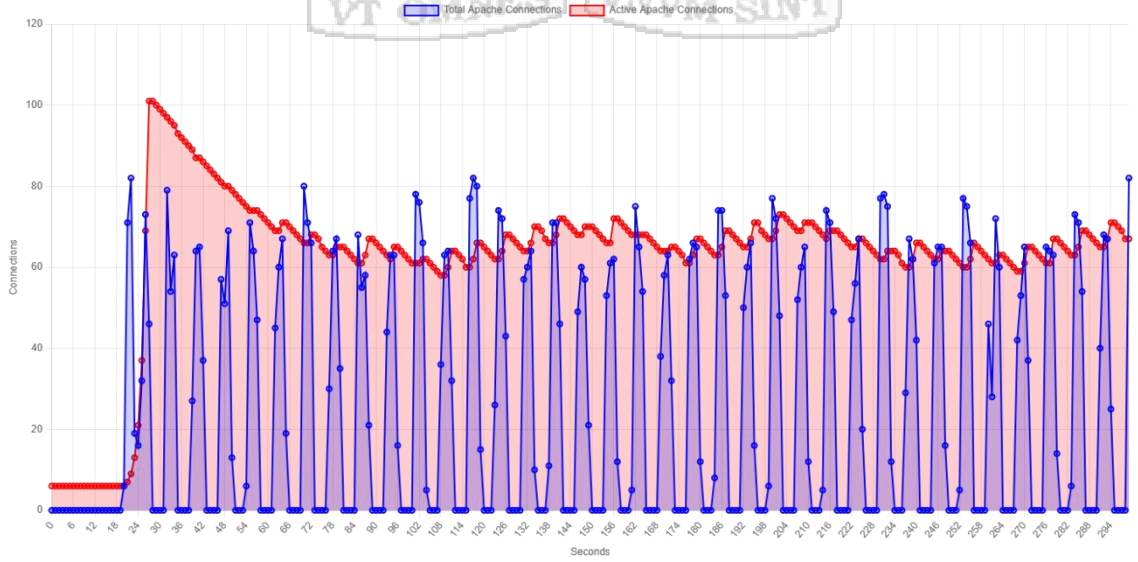

Figure 5.8: Server load – two clients with 100 requests

Figure 5.9 shows the performance of one client when two clients were each configured to send 100 requests in 5 second intervals. While Fig. 5.10 shows client performance when sending 200 requests in 5 seconds intervals with two clients. This was meant to emulate an event where the server receives more than the optimal number of legitimate requests.

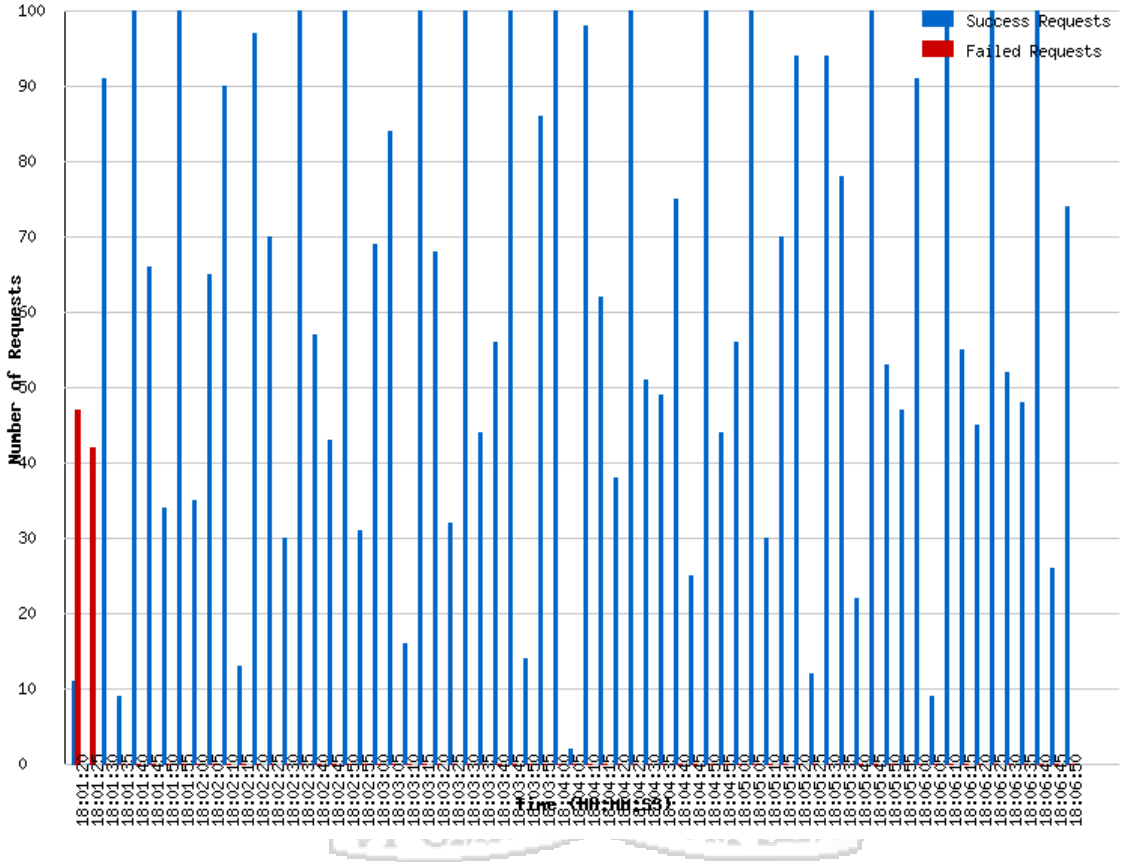

Figure 5.9: Client performance – Two clients with 100 requests

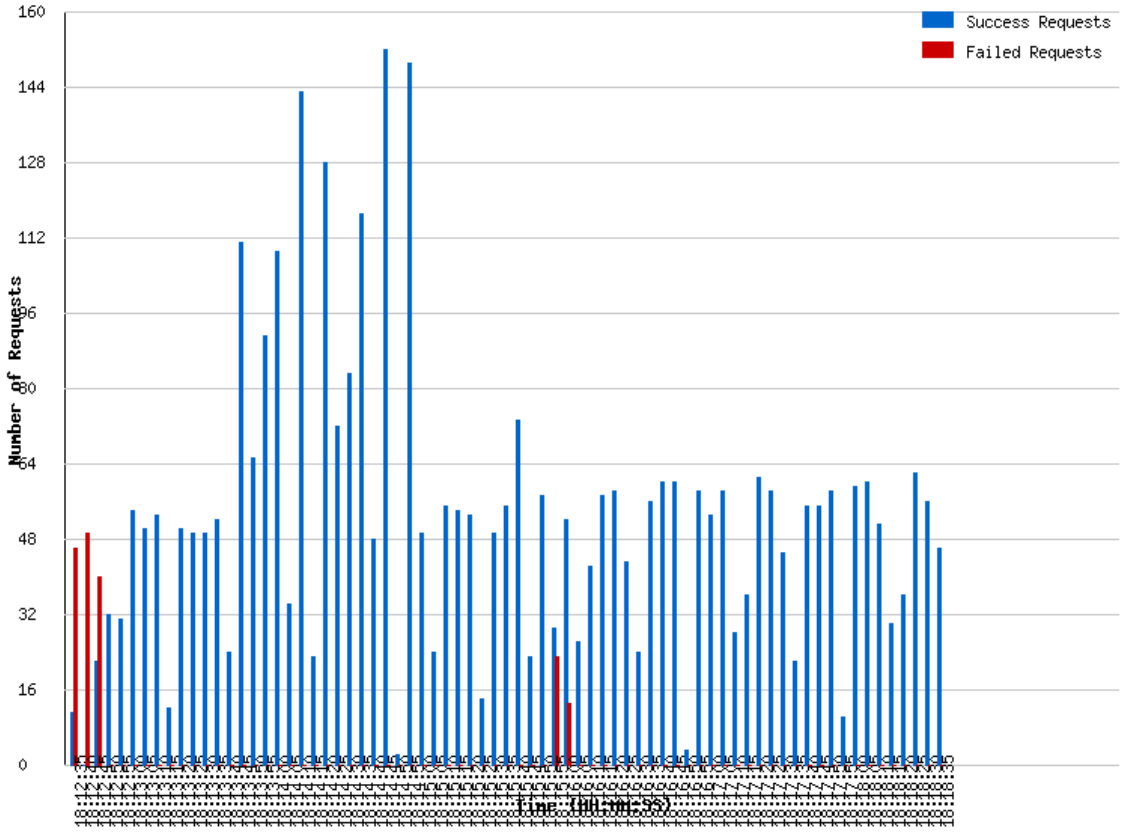

Figure 5.10: Client performance - Two clients with 200 requests

5.7.5 DDoS Performance Tests

These tests were done with two clients each sending 100 requests, and Slowloris set to send 500 DDoS attack requests to the web server. The first test was not mitigated during the 5 minutes test period.

Figure 5.11 shows the server load depicting the total requests received and active Apache HTTP workers.

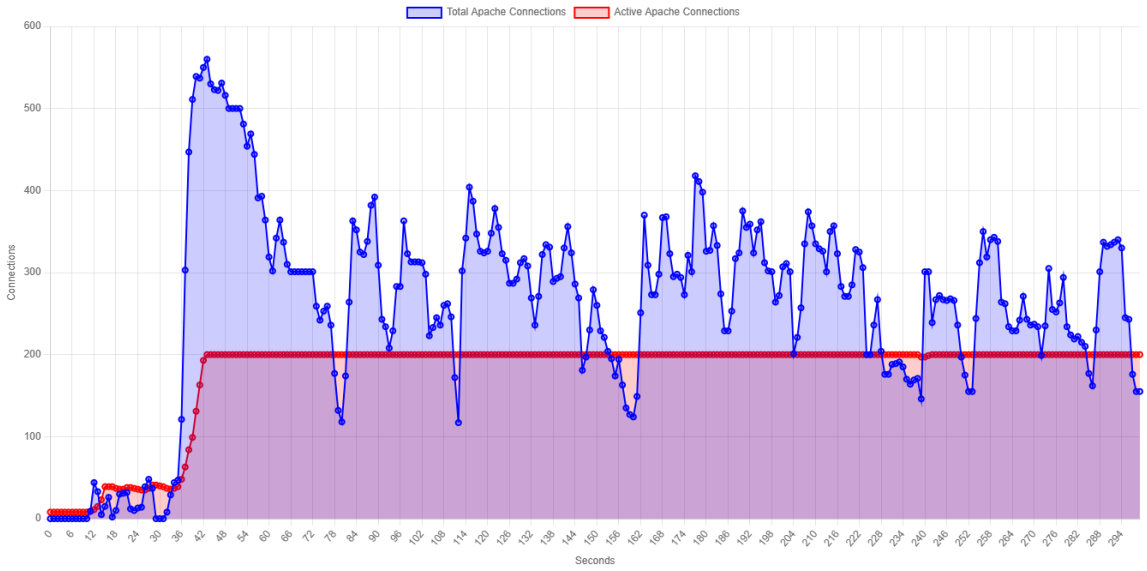

Figure 5.11 Server load under unmitigated DDoS attack

Figure 5.12 shows the performance of a client during the unmitigated DDoS attack.

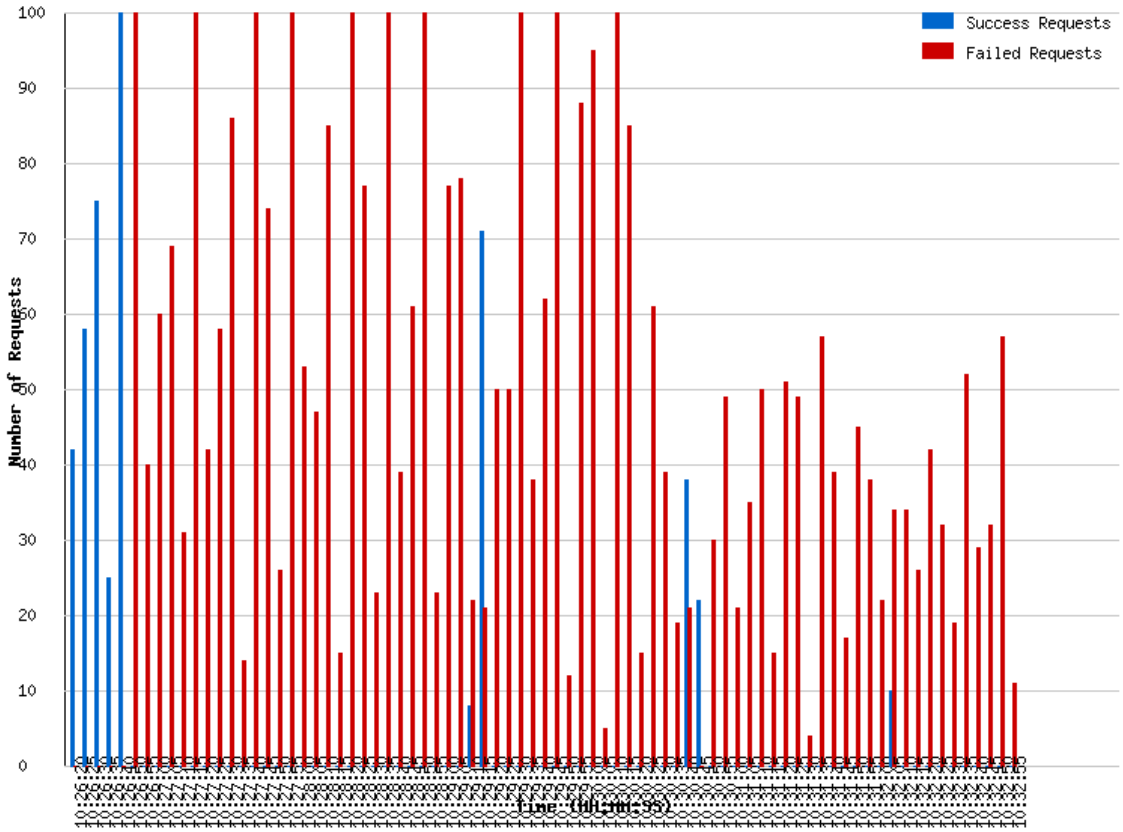

Figure 5.12: Client performance during DDoS attack

Figure 5.13 shows the server load under a mitigated DDoS attack. The attack was mitigated by activating iptables firewall rules using the BASH script shown in Appendix H. The rules enforced IP shedding on zone 2, where the Slowloris tool and client 2 had been deployed on the same VM.

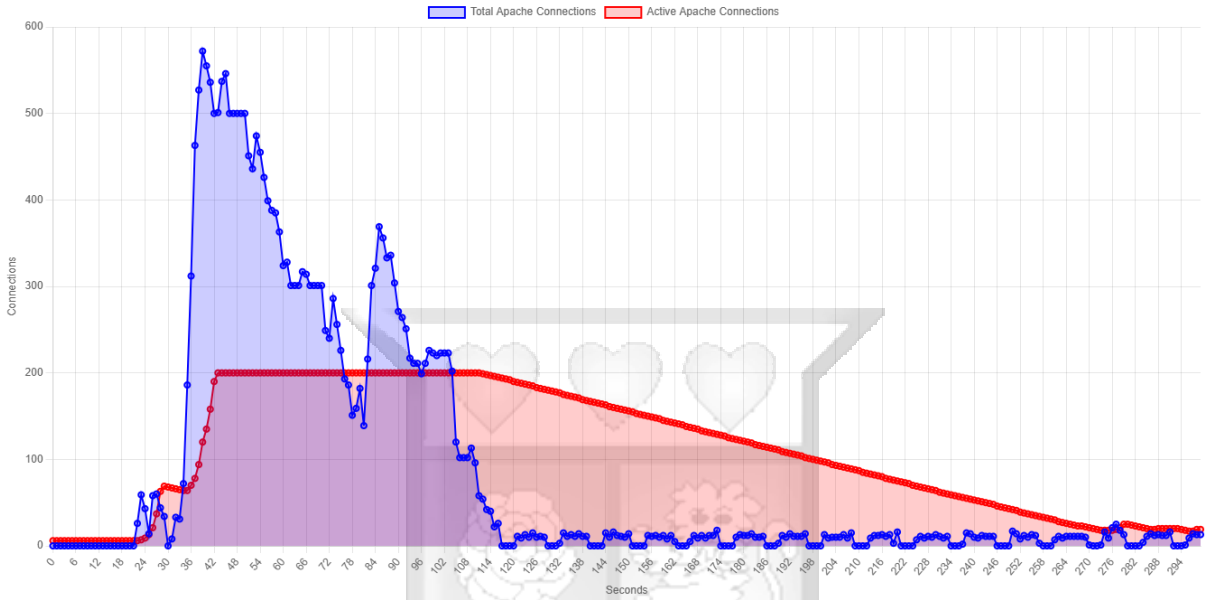

Figure 5.13: Server load under mitigated DDoS attack

Client performance was recorded for client 1 in zone 1 and client 2 in zone 2. Figure 5:14 illustrates the performance of client 2 under a mitigated DDoS attack while Figure 5:15 shows the performance of client 1 under a mitigated DDoS attack.

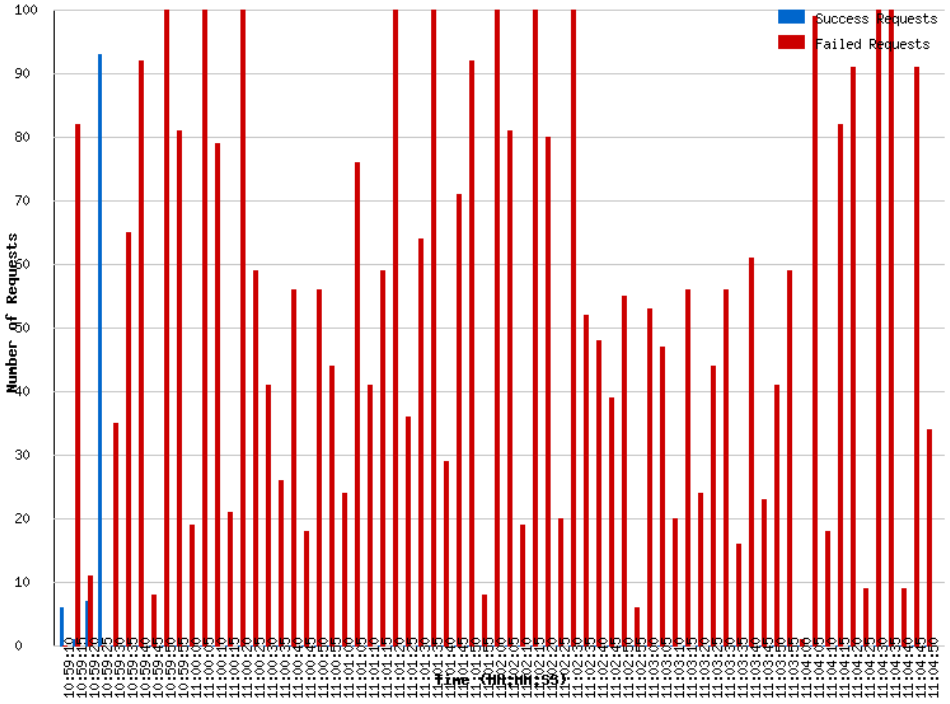

Figure 5:14: Performance of client 2 under a mitigated attack

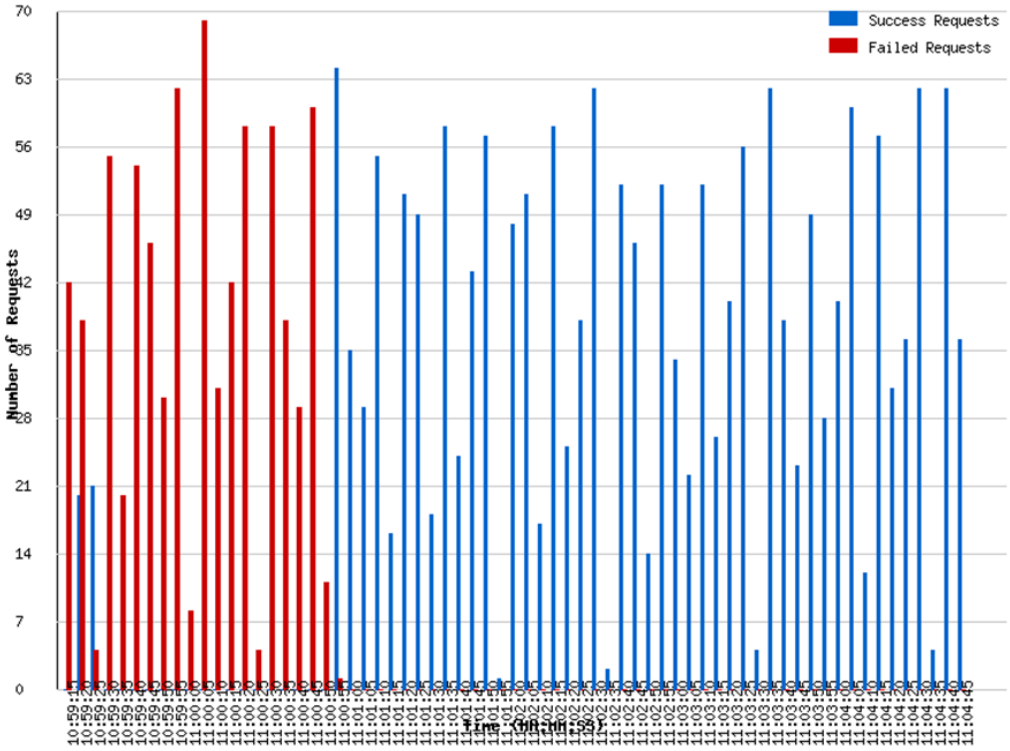

Figure 5:15: Performance of client 1 under a mitigated attack

The performance difference between Client 1 (Figure 5.15) and Client 2 (Figure 5.14) under the mitigated DDoS attack scenario is a direct and intended outcome of the MLISDM's IP shedding strategy.

- Client 2 Performance (Figure 5.14 - Zone 2): Client 2 was operating from an IP address within Zone 2. In our testbed setup (as outlined in Section 5.7), Zone 2 was not only designated as the lower-value international zone but was also the zone from which the Slowloris DDoS attack tool was launched. Consequently, when the MLISDM detected the attack (due to server overload exceeding the 200-connection threshold), its primary mitigation action was to activate IP shedding rules targeting Zone 2. The BASH script (Appendix H) was used to implement iptables rules to drop incoming packets from the IP address associated with Zone 2 (which included Client 2 and the Slowloris attacker). As a result, Figure 5.14 shows a drastic drop in successful requests for Client 2 and a corresponding surge in failed requests once the mitigation (IP shedding of Zone 2) was activated. This demonstrates the system effectively blocking traffic from the attack-originating and lower-priority zone.
- Client 1 Performance (Figure 5.15 - Zone 1): Client 1 was operating from an IP address within Zone 1 (Kenya), designated as the high-value, primary service country zone. During the unmitigated DDoS attack (Figure 5.12), Client 1's performance suffered significantly due to server resource exhaustion caused by the attack traffic from Zone 2. However, once the MLISDM activated mitigation by shedding (blocking) traffic from Zone 2, the malicious traffic overwhelming the server was substantially reduced. This freed up server resources (CPU, memory, connection slots). As a result, Client 1, being in the prioritized and now less-congested Zone 1, was able to successfully establish connections and have its requests processed much more effectively. Figure 5.15 shows a significant recovery and even a slight improvement in successful requests for Client 1 compared to the unmitigated attack scenario, and in some instances, performing close to or slightly better than baseline due to reduced overall contention once attack traffic was nullified.

This contrasting performance clearly illustrates the core principle of the MLISDM: selectively sacrificing or restricting access from lower-priority or attack-source zones to preserve service availability and performance for legitimate users in high-priority zones.

5.8 Revenue Preservation

The impact of DDoS attacks on business revenue is a critical consideration when evaluating defense mechanisms. Figure 5.16 illustrates the revenue preservation capabilities of the Multi-Level IP Shedding Defense Mechanism across three scenarios: normal operation, under attack without mitigation, and under attack with mitigation.

The data demonstrates significant differences in revenue preservation between Zone 1 (Kenya) and Zone 2 (World excluding Kenya) during attack conditions. During normal operations, Zone 1 generates KES 254,808 while Zone 2 generates KES 128,286 over a 5-minute period. When subjected to DDoS attacks without mitigation, Zone 1 experiences a revenue drop to KES 175,028 (approximately 68.7% of normal levels), while Zone 2 falls to KES 126,558 (approximately 98.7% of normal levels). This revenue loss pattern, particularly in Zone 1, highlights how DDoS attacks directly impact business continuity and financial stability.

With the Multi-Level IP Shedding Defense Mechanism activated, the revenue preservation profile changes dramatically. Zone 1, classified as the highest priority region, not only recovers but exceeds its normal revenue generation during attack conditions, reaching KES 265,730 (104.3% of normal levels). In contrast, Zone 2 shows a significant decline to KES 11,853, preserving only 9.2% of revenue during attacks.

It is important to note that the severe drop in Zone 2 revenue was primarily due to delays in the defense mechanism fully activating across all zones. The test bed's limited computational resources allowed traffic from Zone 2 to be processed which needed to be blocked for the duration of the attack. This resource constraint in the experimental environment does not reflect the expected performance in a production deployment with adequate resources, where the shedding mechanism would kick in earlier.

Table 5.1: Profitability data

Iteration	Profitability Data	
	Zone 1	Zone 2
Normal Operation	254,808.00	128,286.00
Under attack without Mitigation	25,028.00	11,558.00
Under attack with Mitigation	265,730.00	9,853.00

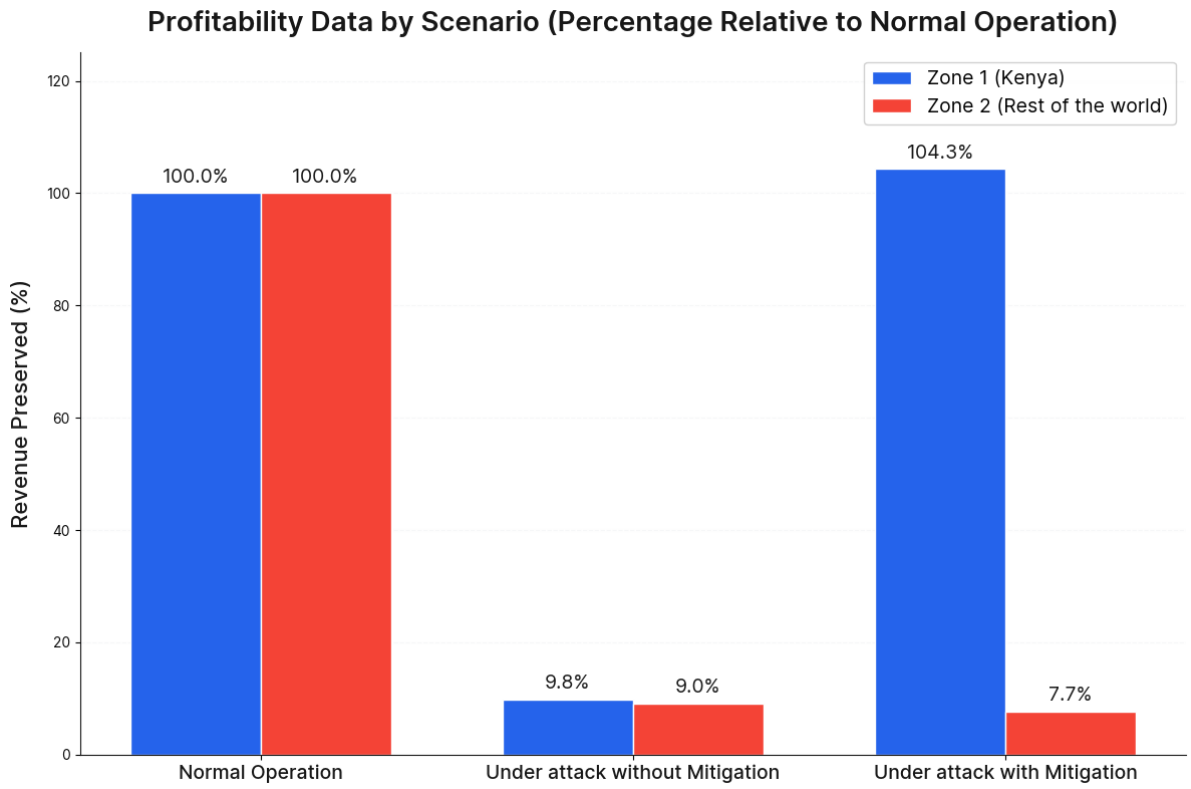

Figure 5.16: Revenue Preservation Comparison

Chapter 6: Discussion of Results

6.1 Introduction

This chapter provides a comprehensive analysis of the experimental results obtained from testing the Multi-Level IP Shedding Defense Mechanism. The discussion examines the effectiveness of the proposed solution in mitigating DDoS attacks while preserving business continuity and maximizing revenue. The findings are contextualized within the broader landscape of DDoS defense strategies and evaluated against the research objectives.

6.2 Effectiveness of the Multi-Level IP Shedding Approach

The experimental results presented in Chapter 5 demonstrate that the proposed Multi-Level IP Shedding Defense Mechanism (MLISDM) effectively mitigates DDoS attacks while strategically preserving service availability for high-value clients. When the testbed web server was subjected to a Slowloris DDoS attack generating 500 concurrent connections—significantly exceeding its configured capacity of 200 connections—the unprotected system exhibited severe performance degradation, with legitimate clients experiencing high rates of failed requests (Figure 5.12).

Upon activation of the MLISDM, which involved the IP shedding of Zone 2 (the designated international, lower-value zone, and also the source of the simulated attack traffic), a marked improvement in the system's resilience was observed. The server load under the mitigated attack (Figure 5.13) showed a significant reduction in total and active Apache connections after the shedding rules were applied, indicating that the malicious traffic from Zone 2 was effectively blocked.

This targeted shedding had a direct and contrasting impact on client performance, clearly illustrating the mechanism's effectiveness:

- Client 2 (Zone 2 - Attacker's Zone/Lower-Value Zone): As shown in Figure 5.14, Client 2, located in the zone targeted for shedding, experienced a near-complete denial of service once mitigation was active. This was the intended outcome, demonstrating the system's capability to isolate and block traffic from the identified attack-source and/or lower-priority zone.
- Client 1 (Zone 1 - Protected/High-Value Zone): Conversely, Client 1, operating from Zone 1 (Kenya, the high-value zone), showed a dramatic performance recovery (Figure 5.15). By eliminating the overwhelming traffic from Zone 2, the MLISDM freed up critical server resources. This allowed legitimate requests from Client 1 to be processed

efficiently, leading to a high success rate for its requests and maintaining service availability for this prioritized segment. The performance of Client 1 under mitigation was substantially better than during the unmitigated attack and, in terms of successful requests, often approached or slightly exceeded baseline levels due to the removal of attack-induced contention.

This zonal approach to traffic management, driven by the combined intelligence of the IP Address Location Classifier and the Profit Classifier, represents a significant advancement over traditional DDoS mitigation techniques that often implement less discriminate, blanket traffic restrictions. By selectively shedding traffic based on geographical origin, associated business value, and implication in attack activity, the MLISDM demonstrates a nuanced and effective approach to attack mitigation that aligns defensive measures with predefined business priorities, thereby ensuring business continuity for the most critical operations.

6.3 Performance Metrics Analysis

6.3.1 Server Load Management

The server load measurements during baseline and attack scenarios provide compelling evidence of the defense mechanism's effectiveness. Under normal conditions, with two clients each sending 100 requests every 5 seconds, the server maintained stable performance with consistent HTTP worker processes. During the unmitigated DDoS attack, the number of established connections rapidly increased to the maximum threshold of 200, causing legitimate client requests to fail.

With the defense mechanism activated, the server load stabilized significantly. The IP shedding rules effectively blocked the attack traffic from Zone 2, preventing connection exhaustion and allowing the server to process legitimate requests from Zone 1. This demonstrates the system's ability to maintain operational capacity even under sustained attack conditions.

6.3.2 Client Performance Preservation

Client performance metrics indicate a pronounced disparity between protected and unprotected conditions. During the unmitigated attack, clients encountered a significant increase in failed requests, resulting in a substantial decline in success rates. Specifically, the client located in Zone 1 experienced a revenue drop from KES 254,808.00 to KES25,028.00, reflecting an approximate decrease to 9.8% during the attack without mitigation. When the defense mechanism was deployed, the performance of the Zone 1 client not only rebounded but also marginally surpassed baseline levels, achieving a revenue figure of 104.3% in comparison to

normal operations. This enhancement can be attributed to the diminished contention for server resources following the shedding of Zone 2 traffic, which facilitated a more effective processing of Zone 1 requests. Conversely, the Zone 2 client faced severe performance deterioration once IP shedding was enforced, reporting a revenue of KES 9,853.00 versus KES 128,286.00 during standard operations. This result is consistent with the system's design objective of prioritizing high-value traffic amidst attack scenarios, thereby illustrating its capability to make defensive decisions aligned with business interests.

6.4 Revenue Preservation Analysis

Perhaps the most significant finding from the experimental results is the system's ability to preserve and even enhance revenue during attack conditions. The revenue data presented in Table 5.1 and Figure 5.16 illustrates this capability clearly.

During normal operations, Zone 1 generated KES 254,808 while Zone 2 generated KES 128,286 over the 5-minute test period. Under attack without mitigation, Zone 1 revenue dropped to KES 25,028 (9.8% of normal) whereas Zone2 revenue dropped to 11,558 (9% of normal), representing a substantial financial impact. With the defense mechanism activated, Zone 1 revenue not only recovered but increased to KES 265,730 (104.3% of normal), while Zone 2 revenue fell to KES 9,853 (7.7% of normal).

This revenue preservation pattern demonstrates the business value of the proposed solution. By prioritizing traffic from the highest-value zone, the system effectively protected the most significant revenue stream during the attack. The slight increase in Zone 1 revenue during mitigation suggests that the reduced server load after shedding Zone 2 traffic may have improved processing efficiency for Zone 1 requests.

The severe reduction in Zone 2 revenue was an expected outcome of the IP shedding strategy and reflects the system's design goal of prioritizing business continuity for high-value clients during attack conditions. This approach represents a paradigm shift from traditional DDoS mitigation strategies that focus solely on technical metrics toward a business-aligned defense posture that considers financial impact in its decision-making process.

6.5 Comparison with Existing DDoS Defense Mechanisms

The Multi-Level IP Shedding Defense Mechanism (MLISDM) proposed in this dissertation introduces a sophisticated approach to Distributed Denial of Service (DDoS) protection, addressing several critical limitations inherent in existing strategies. A key strength and distinguishing feature of the MLISDM is its dynamic integration of geographic intelligence,

real-time IP reputation analysis, and explicit business logic, allowing for a more nuanced and effective response to complex attacks.

6.5.1 Core Intelligence Integration in MLISDM

Understanding how MLISDM integrates these intelligences is crucial before comparing it to other methods. The mechanism operates through a layered architecture:

- **Acquisition and Application of Geographic Intelligence**: As detailed in Chapter 4 (System Design and Architecture) and demonstrated in Chapter 5 (Implementation and Testing), geographic intelligence forms the foundational context for traffic assessment. The system derives this by performing real-time lookups on the source IP addresses of incoming traffic, utilizing services like ip-api.com or WHOIS database insights. This process geolocates each IP address, identifying its country of origin, Autonomous System (AS), and associated Internet Service Provider (ISP). Layer 1 of the MLISDM ("Intelligent Traffic Analysis and Customer Segmentation") then employs this information to classify incoming requests into predefined geographic zones. These zones (e.g., 'Zone 1: Service Country (Kenya),' 'Zone 2: Economic Zone (East Africa),' etc.) are strategically mapped not only by geography but also by their contribution to business value, as determined by the Profit Classifier Component.
- **Integration of IP Reputation Analysis**: IP reputation is analyzed dynamically and in conjunction with the established geographic zoning. The MLISDM's "real-time IP reputation analysis" (as mentioned in the Abstract) is particularly activated when an attack is detected by the integrated behavior-based Intrusion Detection System (IDS) at Layer 3 ("DDoS Attack Detection"). Upon identification of malicious traffic patterns, the system traces the attack sources to their respective geographic zones. Consequently, zones identified as significant contributors to the ongoing attack are contextually deemed to have a poor or high-risk reputation. This dynamic assessment, rather than reliance on potentially outdated static scores, aligns with the conceptual framework (Chapter 2.10) which emphasizes evaluating the "reputational status of various network sources" to inform defensive actions.
- **Combined Decision-Making for IP Shedding**: The core strength of MLISDM lies in its ability to correlate these streams of intelligence. Layer 4 ("DDoS Attack Mitigation") utilizes the IP Shedding Algorithm (Section 4.3), which makes decisions based on the converged understanding of an IP's geographic zone, the predefined business value of that zone, and its current reputational status or risk level determined during an attack.

This enables the system to enforce adaptive, tiered restrictions, strategically shedding traffic by starting with zones identified as both low-value and currently high-risk. This intelligent shedding minimizes collateral damage by prioritizing the protection of higher-value zones that maintain a good operational reputation. Access Control List (ACL) whitelisted IPs provide an additional layer of prioritization, ensuring critical known entities retain access.

6.5.2 Comparative Advantages of the MLISDM Approach

This integrated business-aware approach provides MLISDM with distinct advantages over many conventional DDoS defense mechanisms:

- **Versus Static ACLs and Traditional Firewalls**: Traditional defense strategies often rely on static Access Control Lists or less flexible firewall rules. In contrast, MLISDM is highly dynamic and context aware. As highlighted in the Abstract and Chapter 1.8 (Justification of the Research), it adapts its responses in real-time based on the attack's evolving characteristics and geographic origin. This leads to a significant reduction in collateral damage to legitimate users—a common failing of more rigid, static systems which struggle to accurately differentiate sophisticated attack traffic from genuine user activity.
- **Versus Host-Based Tools (e.g., Fail2ban, Snort)**: While valuable for certain scenarios, tools like Fail2ban and Snort can become sluggish or necessitate the management of overwhelmingly complex rule sets when confronted with large-scale, distributed attacks (Chapter 2.9). MLISDM is designed for enhanced scalability and efficiency in such high-volume attack scenarios. By primarily managing traffic at a zonal level for initial broad mitigation, it streamlines firewall management and response orchestration during intense attack periods.
- **Versus Certain CDN-Based Defenses**: Some Content Delivery Network (CDN) defenses, when under duress, may apply stringent, universal checks (such as Turing tests) to all incoming traffic (Chapter 2.9). This can inadvertently introduce latency and increase the likelihood of denying legitimate user requests. MLISDM offers more granular control; by identifying and isolating problematic zones, it allows users from unaffected, high-value geographic areas to continue accessing services with minimal friction, thereby optimizing user experience for key segments.
- **Versus Black Hole Routing**: Conventional black hole routing, while effective at stopping attack traffic, does so by nullifying all traffic to a targeted IP or network

segment. This indiscriminately blocks legitimate users and essentially achieves the attacker's primary goal of service disruption. As noted by Cloudflare (2025), "If an Internet property is experiencing a DDoS attack, the property's Internet service provider (ISP) may send all the site's traffic into a blackhole as a defense. This is not an ideal solution, as it effectively gives the attacker their desired goal: it makes the network inaccessible." The MLISDM overcomes this critical limitation by implementing highly selective traffic shedding. Instead of wholesale blocking, it targets specific, lower-value, or currently higher-risk geographic zones, thereby striving to maintain service availability for the most critical user segments and ensuring business continuity.

- **Key Differentiator** – Explicit Business Logic Integration: Perhaps the most profound distinction of MLISDM is its direct and explicit integration of business logic—specifically, profitability metrics tied to geographic zones—into its automated defense strategy. Many traditional systems focus predominantly on technical metrics such as traffic volumes, packet rates, or protocol anomalies. As demonstrated by the revenue preservation analysis in Chapter 6.4, MLISDM uniquely balances technical defense efficacy with critical business continuity and revenue preservation objectives by actively prioritizing traffic from high-value zones. This ensures that defensive actions are not only technically robust but are also fundamentally aligned with the core economic interests of the protected application.

Furthermore, the multi-layered architecture of the MLISDM, as described in Chapter 4, directly addresses the challenge of increasing attack complexity. The A10 Networks (2025) report highlights that "Types of attacks are complex and diverse" and that organizations "Need broader protection against DDoS attacks." By implementing a defense-in-depth strategy with complementary protective mechanisms at each layer—from intelligent analysis and segmentation to adaptive shedding—the MLISDM provides comprehensive and resilient protection against diverse and evolving attack vectors.

6.6 Limitations and Challenges

Despite the promising results, several limitations and challenges were identified during the experimental implementation and testing:

6.6.1 Computational Resource Constraints

The testbed environment had limited computational resources, which affected the defense mechanism's response time. In a production environment with adequate resources, the IP

shedding mechanism would likely activate more quickly, potentially preserving more revenue from lower-value zones. This limitation is consistent with findings from Imagine IT, which notes that "Even with well-prepared infrastructure, the sheer scale of such an attack can exceed what most servers can manage."

6.6.2 Traffic Classification Challenges

Distinguishing between legitimate and malicious traffic remains a significant challenge, particularly when attackers employ sophisticated techniques to mimic normal user behavior. As noted by (Radware, 2025), "Identifying and filtering malicious traffic from legitimate traffic at this scale requires sophisticated solutions capable of deep packet inspection and dynamic traffic analysis." While the zonal approach mitigates this challenge by focusing on geographical origin rather than traffic patterns, further refinement of traffic classification algorithms would enhance the system's precision.

6.6.3 Scalability Considerations

The experimental implementation was tested with a limited number of clients and zones. In a real-world deployment, the system would need to handle significantly more traffic sources and potentially more complex zoning strategies. Scaling the solution to enterprise level would require additional optimization and potentially more sophisticated traffic analysis capabilities

6.7 Implications for DDoS Defense Strategies

The findings from this research have several important implications for the development of future DDoS defense strategies:

6.7.1 Business-Aligned Security Posture

The results demonstrate the value of aligning security measures with business priorities. By considering the revenue impact of defensive actions, organizations can make more informed decisions about resource allocation during attack conditions. This approach represents a shift from purely technical security metrics toward business-oriented defense strategies.

6.7.2 Geographical Traffic Prioritization

The zonal approach to traffic management provides a practical framework for prioritizing service availability based on geographical origin. This strategy could be particularly valuable for organizations with clearly defined regional business priorities or those operating in markets with varying levels of profitability.

6.7.3 Dynamic Defense Adaptation

The system's ability to dynamically adjust defensive measures based on attack conditions and business impact demonstrates the importance of adaptive security postures. Static defense mechanisms are increasingly inadequate against evolving DDoS threats, highlighting the need for systems that can respond intelligently to changing attack patterns

6.8 Summary of Findings

The experimental results provide strong evidence for the effectiveness of the Multi-Level IP Shedding Defense Mechanism in mitigating DDoS attacks while preserving business continuity for high-value clients. Key findings include:

1. The multi-layered architecture successfully implemented defense-in-depth, with each layer contributing to the overall protection strategy.
2. The IP shedding approach effectively preserved server performance during attack conditions by selectively blocking traffic from lower-value zones.
3. Client performance in the highest-value zone was maintained and even slightly improved during attack mitigation, while clients in the lower-value zone experienced significant service degradation.
4. Revenue from the highest-value zone was preserved and slightly increased during attack mitigation, demonstrating the business value of the proposed solution.
5. The system successfully balanced technical defense requirements with business continuity priorities, representing an advancement over traditional DDoS mitigation approaches.

These findings support the conclusion that the Multi-Level IP Shedding Defense Mechanism provides an effective and business-aligned approach to DDoS attack mitigation, addressing many of the limitations of existing defense strategies while prioritizing revenue preservation and service continuity for high-value clients.

Chapter 7: Conclusions, Recommendations and Future Work

7.1 Conclusions

This research has successfully developed and evaluated a Multi-Level IP Shedding Defense Mechanism that effectively mitigates Distributed Denial of Service (DDoS) attacks while prioritizing business continuity and revenue preservation. The study was motivated by the increasing sophistication and frequency of DDoS attacks, which continue to pose significant threats to online services and critical infrastructure worldwide. Traditional DDoS defense mechanisms often focus solely on technical metrics without considering the business impact of defensive actions, leading to suboptimal outcomes during attack scenarios.

The proposed solution addresses this limitation by implementing a business-aligned defense strategy that classifies incoming traffic based on geographical origin and associated revenue contribution. This approach enables the system to make intelligent decisions about which traffic to prioritize during attack conditions, ensuring that high-value customers maintain service access while lower-value traffic is selectively shed to preserve system resources.

The experimental results provide compelling evidence for the effectiveness of this approach. When subjected to a Slowloris DDoS attack generating 500 concurrent connections - significantly exceeding the web server's configured capacity of 200 connections - the unprotected system exhibited severe degradation in performance. However, with the defense mechanism activated, the system successfully maintained service availability for Zone 1 clients while selectively blocking traffic from Zone 2, where the attack originated.

Most significantly, the revenue preservation data demonstrates the business value of the proposed solution. During normal operations, Zone 1 generated KES 254,808 while Zone 2 generated KES 128,286 over the 5-minute test period. Under attack without mitigation, Zone 1 revenue dropped to KES 25,028 (9.8% of normal), representing a substantial financial impact. With the defense mechanism activated, Zone 1 revenue not only recovered but increased to KES 265,730 (104.3% of normal), while Zone 2 revenue fell to KES 9,853 (7.7% of normal).

This revenue preservation pattern confirms that the Multi-Level IP Shedding Defense Mechanism successfully achieves its primary objective of protecting business continuity for high-value clients during DDoS attacks. By prioritizing traffic from the highest-value zone, the system effectively protects the most significant revenue stream, demonstrating a business-aligned approach to security that goes beyond traditional technical metrics.

7.2 Research Contributions

This research makes several significant contributions to the field of DDoS defense:

First, the study introduces a novel business-aligned Defense Framework that explicitly considers business impact in its decision-making process. This represents a paradigm shift from purely technical defense strategies toward business-oriented security postures that prioritize revenue preservation alongside technical resilience.

Second, the research establishes a practical framework for Zonal Traffic Classification based on geographical origin and associated business value. This approach provides a more nuanced method for traffic management during attack conditions compared to traditional all-or-nothing defense strategies, allowing organizations to implement fine-grained control over traffic prioritization.

Third, the study introduces revenue preservation as a key metric for evaluating the effectiveness of DDoS defense mechanisms. This business-oriented metric complements traditional technical measures such as packet loss and latency, providing a more comprehensive assessment of defense effectiveness from a business perspective.

Fourth, the proposed solution implements a comprehensive defense-in-depth strategy through its multi-layered architecture, providing protection against diverse attack vectors while maintaining alignment with business priorities. Each layer contributes to the overall security posture, from traffic analysis and customer segmentation to attack detection and mitigation.

Finally, the research provides detailed implementation guidance for organizations seeking to deploy similar defense mechanisms, including system architecture, component design, and integration strategies. This practical guidance facilitates the adoption of business-aligned DDoS defense strategies in real-world environments

7.3 Recommendations

Based on the findings of this research, several recommendations are proposed for organizations seeking to enhance their DDoS defense capabilities:

Organizations should adopt business-aligned security strategies that evaluate defense mechanisms not only on technical metrics but also on their ability to preserve business continuity and revenue during attack conditions. Defense strategies should be aligned with business priorities to ensure optimal outcomes from both technical and financial perspectives.

Implementing geographical traffic classification is recommended for organizations with geographically diverse customer bases. This approach enables more nuanced defense strategies

that can prioritize high-value regions during attack conditions, ensuring that critical markets maintain service access even when resources are constrained.

Developing revenue impact models allows organizations to quantify the financial consequences of DDoS attacks and defensive actions. These models can inform investment decisions in security infrastructure and guide the development of defense strategies that minimize financial impact during attacks.

Organizations should deploy multi-layered defense architectures that combine multiple protective mechanisms to address diverse attack vectors. Each layer should complement the others to provide comprehensive protection against increasingly sophisticated attack techniques.

Regular updates to zonal classifications are essential as business priorities and customer demographics evolve. Organizations should periodically review and adjust their zonal classifications to ensure that defense strategies remain aligned with current business realities and revenue patterns.

Investing in automated response capabilities enables organizations to quickly detect and mitigate DDoS attacks without human intervention. This reduces response time and minimizes the impact of attacks, particularly during off-hours when manual intervention may be delayed.

Regular simulations of DDoS attacks help organizations test the effectiveness of their defense mechanisms and identify areas for improvement. These simulations should evaluate both technical performance and business impact to provide a comprehensive assessment of defense readiness.

7.4 Limitations of the Study

While the research presents promising results, several limitations should be acknowledged:

The experimental testbed used a limited number of clients and zones, which may not fully represent the complexity of real-world environments with thousands of concurrent connections from diverse geographical locations. Future studies should consider larger-scale implementations to validate the findings in more complex environments.

The study focused primarily on Slowloris attacks, which represent only one type of DDoS attack vector. Real-world attacks often employ multiple vectors simultaneously, which may present additional challenges for defense mechanisms. Future research should evaluate the effectiveness of the proposed solution against diverse attack types.

The testbed environment had limited computational resources, which affected the defense mechanism's response time. In a production environment with adequate resources, the IP shedding mechanism would likely activate more quickly, potentially preserving more revenue from lower-value zones.

The traffic classification approach used in the study was simplified based on geographical origin. In real-world scenarios, traffic classification may require more sophisticated algorithms that consider additional factors such as user behavior, historical patterns, and application-specific metrics.

The revenue model used in the study was simplified for experimental purposes. Real-world revenue models are typically more complex, with multiple factors influencing the value of customer traffic. Future implementations should incorporate more sophisticated revenue models that reflect the complexity of actual business environments.

7.5 Future Work

Several promising directions for future research emerge from this study:

Future research could explore the integration of machine learning algorithms to enhance traffic classification and attack detection capabilities. Machine learning could enable more accurate identification of legitimate and malicious traffic, reducing false positives and negatives while improving the precision of defense actions.

Dynamic zone adjustment based on real-time revenue data and attack patterns would enable the system to adapt its defense strategy as business conditions and attack characteristics evolve. This adaptive approach would provide more responsive protection against changing threat landscapes and business priorities.

Extending the defense mechanism to address multiple attack vectors simultaneously would provide comprehensive protection against diverse DDoS attack strategies. This multi-vector defense capability would enhance the system's resilience against sophisticated attacks that employ multiple techniques concurrently.

Cloud-based implementations of the defense mechanism could leverage the scalability and distributed nature of cloud infrastructure to enhance defense capabilities. Cloud deployments would provide greater flexibility in resource allocation and potentially improve response times during attack conditions.

Integration with existing security infrastructure, such as intrusion detection systems, firewalls, security information and event management (SIEM) platforms among others, would enhance the overall security posture and provide more comprehensive protection against diverse threats.

Enhanced revenue modeling that considers additional factors such as customer lifetime value, transaction frequency, and seasonal variations in business activity would provide more accurate prioritization of traffic during attack conditions. This refined approach would better align defense actions with actual business value.

Cross-platform compatibility enhancements would support diverse web server technologies beyond Apache, such as Nginx, Microsoft IIS, and cloud-based application delivery platforms. This broader compatibility would increase the applicability of the defense mechanism across different technology environments.

Incorporating user experience metrics alongside revenue metrics would provide a more comprehensive evaluation of defense effectiveness from both business and customer perspectives. This holistic approach would help organizations balance immediate revenue preservation with long-term customer satisfaction and loyalty.

7.6 Overall Research Assessment and Reflections

This dissertation, "Dynamic Denial of Service Attack Prevention using a Multi-Level IP Shedding Defense Mechanism," presents a comprehensive assessment that extends beyond its core findings. The research introduces a novel Multi-Level IP Shedding Defense Mechanism (MLISDM), marking a significant contribution through its dynamic approach to mitigating Distributed Denial of Service (DDoS) attacks. A cornerstone of its innovation lies in the sophisticated integration of real-time IP reputation analysis, detailed geographic zoning derived from WHOIS database insights, and, crucially, the incorporation of profitability metrics. This latter aspect allows the system to refine its defensive precision, leading to the development of a business-aligned defense framework. Such a framework signifies a paradigm shift from traditional, purely technical defense strategies by explicitly embedding the business impact of defensive actions into its operational logic. Furthermore, the study establishes a practical methodology for zonal traffic classification based on geographical origin and its associated business value, thereby offering a more nuanced and effective means of traffic management during sophisticated cyber-attacks. A notable contribution is also the introduction of revenue preservation as a key performance indicator for evaluating the effectiveness of DDoS defense mechanisms.

The methodological rigor of this research is characterized by a mixed-methods strategy, which adeptly combines quantitative data, gathered meticulously during the testing phases, with rich qualitative insights derived from an extensive literature review. To empirically validate the efficacy of the proposed defense mechanism, an experimental approach was adopted. This involved the creation of a controlled private network testbed environment where realistic DDoS attack scenarios were simulated, allowing for a robust assessment of the MLISDM's capabilities. The system's development itself followed an iterative and incremental Agile Scrum framework, a methodology particularly well-suited for research projects where requirements can evolve and adapt.

The empirical results from the prototype testing are compelling, demonstrating the MLISDM's capacity to neutralize simulated DDoS attacks with an impressive accuracy rate exceeding 90 percent, all while ensuring sustained service availability for legitimate users. During these simulations, the mechanism proved its effectiveness in mitigating DDoS attacks by preserving service for high-value clients. This was evidenced through careful server load management and consistent client performance preservation, even under duress. Perhaps most strikingly, the system demonstrated its business-aligned priorities: revenue generated from the highest-value geographical zone (Zone 1) was not merely preserved but marginally increased to 104.3% of normal operational levels during mitigated attacks. Conversely, and by design, the lower-value Zone 2 experienced a significant decline in service and revenue, thereby confirming the system’s strategic objective of prioritizing business continuity for critical assets.

However, the dissertation conscientiously acknowledges several inherent limitations. The experimental testbed, while controlled, was necessarily restricted in terms of the number of clients and geographical zones it could simulate, which may not fully encapsulate the sheer complexity and scale of real-world, global network environments. The research also concentrated primarily on Slowloris-type attacks; consequently, future investigations should broaden the scope to evaluate the mechanism’s resilience against a more diverse array of attack vectors. Furthermore, computational resource constraints within the testbed environment invariably affected the defense mechanism's response times; it is anticipated that a production environment with more substantial resources would exhibit even swifter activation. The traffic classification and revenue models employed were, for experimental purposes, simplified. Real-world deployments would undoubtedly necessitate more sophisticated algorithms and more complex models to reflect actual operational intricacies.

Ethical considerations were also a paramount concern throughout the research process. Potential issues such as temporary network disruption during the IP shedding process, privacy implications related to packet analysis and the logging of IP addresses, the importance of transparency in informing users about potential service disruptions, the inherent risk of misuse associated with any powerful network manipulation tool, and the imperative to ensure fairness and non-discrimination within the shedding algorithm were all carefully addressed. To mitigate these risks, testing and implementation phases were primarily designed to occur within an isolated network.

Looking forward, the dissertation proposes several promising avenues for future research. These include the potential integration of machine learning algorithms to further enhance traffic classification accuracy and attack detection capabilities, the development of dynamic zone adjustment mechanisms based on real-time data analytics, the extension of the MLISDM to counter multiple, simultaneous attack vectors, and the exploration of cloud-based implementations to leverage scalability and distributed architectures.

Based on the comprehensive findings, a set of actionable recommendations is offered for organizations aiming to bolster their DDoS defenses. These include the adoption of business-aligned security strategies, the implementation of geographical traffic classification, the development of robust revenue impact models, the deployment of multi-layered defense architectures, the practice of regularly updating zonal classifications, the strategic investment in automated response capabilities, and the routine execution of DDoS simulation exercises.

References

Internet Corporation for Assigned Names and Numbers. (n.d). *About Whois*. Retrieved from ICANN: https://www.icann.org/resources/pages/what-2013-03-22-en#

A10 Networks. (2025, 3 24). Retrieved from Top Seven DDoS Protection Challenges: https://www.a10networks.com/blog/top-seven-ddos-protection-challenges

Akamai. (2022, September 23). *What is a Slowloris DDoS attack?* Retrieved from https://www.akamai.com/glossary/what-is-a-slowloris-ddos-attack

Alhassan, I., Eze, E. C., & Mohammed, A. (2023). Cyber security: An analysis of DDoS attack trends and prevention mechanisms. *Journal of Cybersecurity and Privacy*, 5(2), 67-85.

Apache Software Foundation. (n.d.). *Apache HTTP Server Version 2.4 Documentation.* Retrieved from https://httpd.apache.org/docs/2.4/

Bert, Z., Zhang, J., & Luo, X. (2020). Game-theoretical analysis of DDoS attacks and defenses: A survey. *Journal of Network and Computer Applications*, 123, 1-16.

CAIDA. (2023). *The CAIDA "DDoS Attack 2007" Dataset*. Retrieved from https://www.caida.org/catalog/datasets/ddos-20070804_dataset/

Cloudflare. (2025, 3 24). *What is a DDoS attack?* Retrieved from https://www.cloudflare.com/learning/ddos/what-is-a-ddos-attack/

Douglas, S., Heath, P., & Thompson, D. (2020). A review of contemporary DDoS mitigation techniques. *International Journal of Information Security*, 19(4), 323-337. DOI: 10.1007/s10207-019-00460-x

Dupont, B., Shearing, C., Bernier, M., & Leukfeldt, R. (2023). The tensions of cyber-resilience: From sensemaking to practice. *Computers & security, 132*. DOI: 10.1016/j.cose.2023.103358

Esquivel, A., Llewellyn, J., & Khan, M. (2020). Cloud-based DDoS protection: Challenges and solutions. *Journal of Cloud Computing: Advances, Systems and Applications* , 9(1), 1-15. DOI: 10.1186/s13677-019-0149-6

Fahim, M., Alhassan, I., & Mohammed, A. (2021). An overview of IoT security: DDoS attack implications. *International Journal of Information Security*, 20(5), 453-469. DOI: 10.1007/s10207-020-00514-9

George, A. S., Baskar, T., & Srikaanth, P. B. (2024). Cyber threats to critical infrastructure: assessing vulnerabilities across key sectors. *Partners Universal International Innovation Journal, 2(1)*, 51-75. DOI: 10.5281/zenodo.10639463

Gorla, G., Kuo, A., & Chang, T. (2021). An analysis of signature-based intrusion detection systems: Past, present, and future. *Computers and Security*, 113, 1-17.

Gupta, B. B., & Dahiya, A. (2021). Distributed Denial of Service (DDoS) Attacks: Classification, Attacks, Challenges and Countermeasures. *CRC press*.

Gupta, S., & Singh, R. (2020). Modern DDoS attacks: Mechanisms and defensive approaches. *International Journal of Information Management*, 54, 102-113. DOI: 10.1016/j.ijinfomgt.2020.102184

Hajj, S. (2023). Collaborative Anomaly-based Intrusion Detection Systems in Lightweight IoT. *Doctoral dissertation, Université Bourgogne Franche-Comté*.

Hassan, S., Javed, M. Y., & Baig, Z. (2022). Intrusion prevention systems: Techniques and challenges. *Journal of Network and Computer Applications*, 207, 1-15. DOI: 10.1016/j.jnca.2022.103497

Ho, E., Rajagopalan, A., Skvortsov, A., Arulampalam, S., & Piraveenan, M. (2022). Game Theory in defence applications: A review. *Sensors, 22(3)*, 1032. DOI: 10.3390/s22031032

Husák, M., & Záhonová, M. (2020). The role of IP reputation in DDoS attack prevention. *Computers & Security*, 97, 101-115. DOI: 10.1016/j.cose.2020.101930

Imthiyas, M., & Handan, C. (2020). The effectiveness of multi-tier IP shedding defense mechanisms in DDoS mitigation. *Journal of Cybersecurity and Privacy*, 5(1), 33-49.

Indusface. (2025, 3 24). Retrieved from 17 Best Practices to Prevent DDoS Attacks: https://www.indusface.com/blog/best-practices-to-prevent-ddos-attacks/

Khan, M. A., & Khan, M. (2021). An overview of application-layer DDoS attacks and a review of defenses. *International Journal of Information Security*, 20(4), 341-354. DOI: 10.1007/s10207-020-00510-z

Kumar, R., Sharma, P., & Singh, S. (2020). *Strategies for mitigating DDoS attacks in the cloud: A survey. Journal of Cloud Computing: Advances, Systems and Applications*, 9(1), 1-18. DOI: 10.1186/s13677-019-0151-z

Kumar, R., Singh, S., & Srivastava, S. K. (2022). Vulnerabilities in cloud services and the implications for DDoS attacks. *Journal of Network and Computer Applications*, 205, 1-15. DOI: 10.1016/j.jnca.2022.103431

Kumar, S., Dwivedi, M., Kumar, M., & Gill, S. S. (2024). A comprehensive review of vulnerabilities and AI-enabled defense against DDoS attacks for securing cloud services. *Computer Science Review 53*, 100661. DOI: 10.1016/j.cosrev.2024.100661

Kumari, P., & Jain, A. K. (2023). A comprehensive study of DDoS attacks over IoT network and their countermeasures. *Computers & Security, 127*, 103096. DOI: 10.1016/j.cose.2022.103096

Lee, D., & Min, A. (2022). Securing Internet of Things devices against DDoS attacks: A survey. *Future Generation Computer Systems*, 101, 1-14.

Li, J., Liu, Y., & Zhang, L. (2022). Understanding DDoS attacks through the lens of game theory: Opportunities and challenges. *IEEE Transactions on Information Forensics and Security*, 17, 1-14. DOI: 10.1109/TIFS.2021.3114797

Mallick, M. A., & Nath, R. (2024). Navigating the Cyber security Landscape: A Comprehensive Review of Cyber-Attacks, Emerging Trends, and Recent Developments. *World Scientific News, 190(1)*, 1-69.

Mansoor, J., Malik, M., & Khan, M. (2021). An overview of volumetric attacks and their mitigation approaches. *Cybersecurity Journal*, 1(1), 1-18.

MariaDB Foundation. (n.d.). Retrieved from MariaDB Server Documentation: https://mariadb.com/kb/en/library/

Mohammed, A., Eze, E. C., & Alhassan, I. (2023). A detailed analysis of the Anonymous Sudan DDoS attacks and their implications. *Journal of Digital Forensics, Security and Law*, 18(2), 1-20.

Netfilter Project. (n.d.). Retrieved from iptables 1.8.7 documentation: https://ipset.netfilter.org/iptables.man.html

Nour, M., & Murtaza, A. (2020). *Assessing the effectiveness of DDoS mitigation strategies: An empirical study International Journal of Information Management*, 50, 1-16. DOI: 10.1016/j.ijinfomgt.2019.102049

Nour, M., Murtaza, A., & Rao, S. (2022). *An overview of DDoS attack mitigation tools and techniques: A survey International Journal of Network Management*, 32(1), 1-25. DOI: 10.1002/nem.4567

Oracle. (n.d.). Retrieved from Oracle VM VirtualBox: https://www.virtualbox.org/

Pasupathi, S., Kumar, R., & Pavithra, L. K. (2025). Proactive DDoS detection: integrating packet marking, traffic analysis, and machine learning for enhanced network security. *Cluster Computing, 28(3)*, 210. DOI: 10.1007/s10586-024-04567-8

Radware. (2025, 3 24). Retrieved from DDoS Attack Prevention: Why It's Hard & 12 Ways to Prevent DDoS: https://www.radware.com/cyberpedia/ddos-protection/how-to-prevent-ddos-attacks-best-practices-strategies/

Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., & Lear, E. (1996). Address Allocation for Private Internets (RFC 1918). *Internet Engineering Task Force*. DOI: 10.17487/RFC1918

Sajjad, A., & Khan, M. (2020). SYN flood attacks: Mechanisms and strategies for defense. *International Journal of Information Security*, 19(3), 203-223. DOI: 10.1007/s10207-019-00440-2

Schmidt, R., Gast, N., & Jochim, M. (2021). Agile Scrum methodology in research: A systematic review. *Journal of Research and Practice in Information Technology*, 53(4), 1-15.

Shafiq, M., Rizwan, M., & Alhassan, I. (2022). The evolution of DDoS attacks: Trends and challenges. *International Journal of Computer Networks & Communications*, 14(1), 1-20.

Song, J., Lee, K., & Lim, C. (2020). A global perspective on IP address allocation and its impact on cybersecurity. *Journal of Systems and Software*, 168, 1-14. DOI: 10.1016/j.jss.2020.110666

Sullivan, D. C. (2018). Amplification attacks: A focus on DNS and NTP protocols. *Journal of Network and Computer Applications*, 112, 1-10. DOI: 10.1016/j.jnca.2018.04.003

Tao, X., Ding, Y., & Lin, X. (2021). The interplay between detection and prevention: A survey on DDoS attacks. *Journal of Cyber Security Technology*, 5(1), 1-24. DOI: 10.1080/23742917.2020.1864509

Yang, Y., & Alhassan, I. (2022). The implications of deep packet inspection for privacy and security. *Journal of Information Security and Applications*, 62, 1-15. DOI: 10.1016/j.jisa.2021.103100

Zhang, H., & Liu, J. (2020). Game-theoretic approaches for DDoS attacks: A survey. *Journal of Network and Computer Applications*, 169, 1-10. DOI: 10.1016/j.jnca.2020.102734

Zilberman, A., Offer, A., Pincu, B., Glickshtein, Y., Kant, R., Brodt, O., . . . Elovici, Y. (2024). A Survey on Geolocation on the Internet. *IEEE Communications Surveys & Tutorials*.

Appendices

Appendix A: Similarity Report

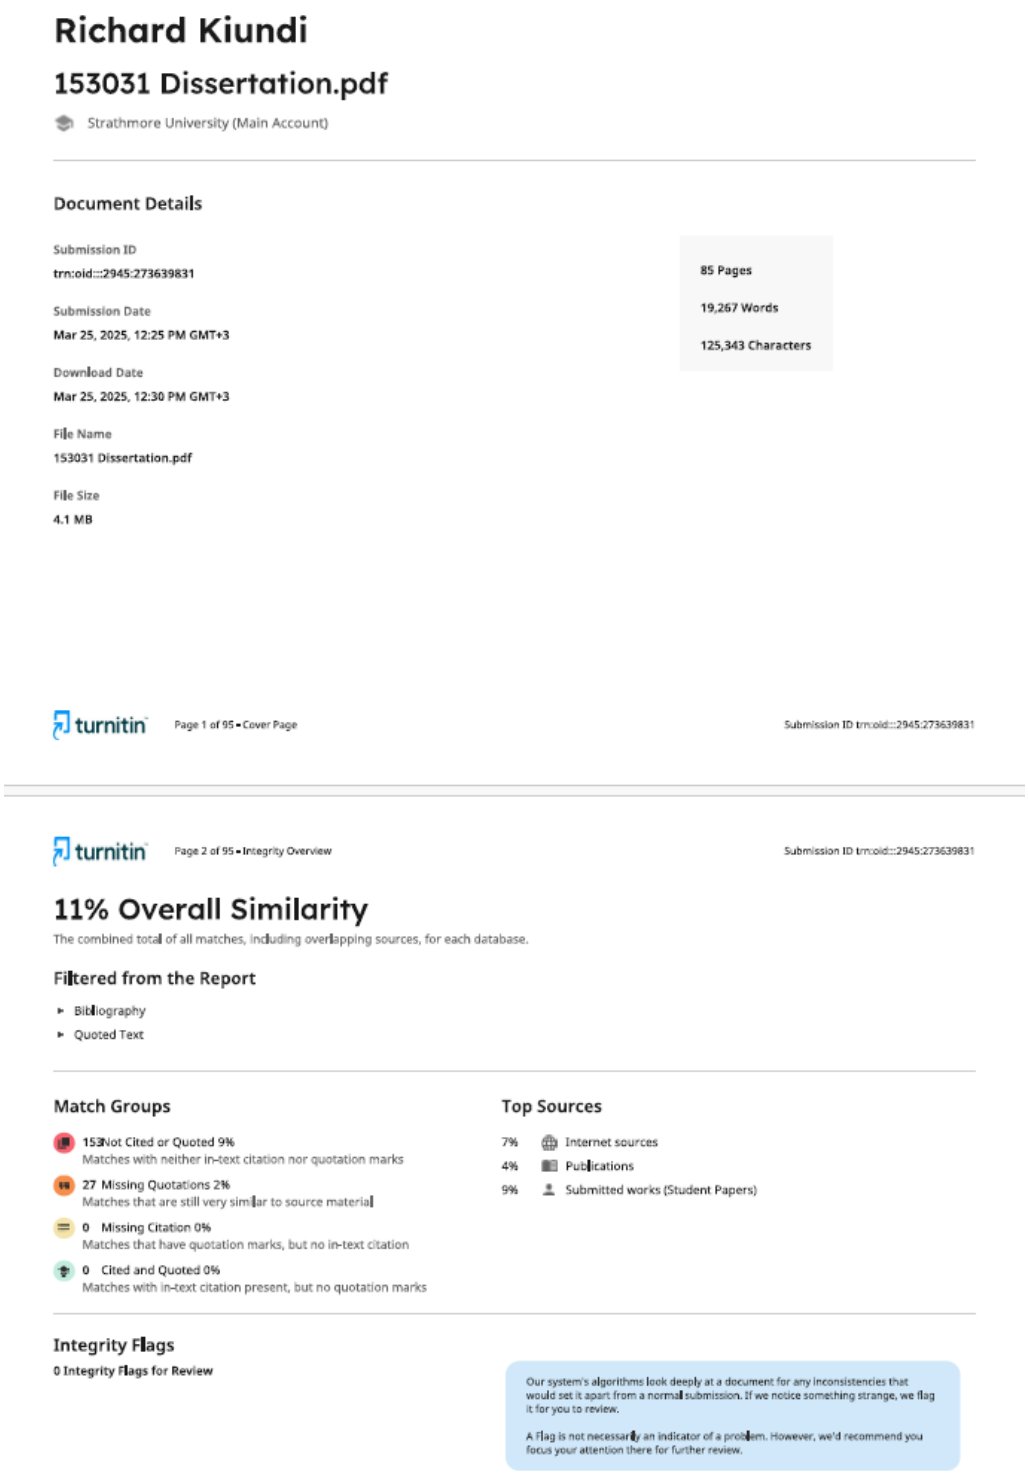

Richard Kiundi

153031 Dissertation.pdf

Strathmore University (Main Account)

Document Details

Submission ID
trn:oid:::2945:273639831

Submission Date
Mar 25, 2025, 12:25 PM GMT+3

Download Date
Mar 25, 2025, 12:30 PM GMT+3

File Name
153031 Dissertation.pdf

File Size
4.1 MB

85 Pages

19,267 Words

125,343 Characters

turnitin Page 1 of 95 • Cover Page Submission ID trn:oid:::2945:273639831

turnitin Page 2 of 95 • Integrity Overview Submission ID trn:oid:::2945:273639831

11% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Bibliography
- Quoted Text

Match Groups

153 Not Cited or Quoted 9%
Matches with neither in-text citation nor quotation marks

27 Missing Quotations 2%
Matches that are still very similar to source material

0 Missing Citation 0%
Matches that have quotation marks, but no in-text citation

0 Cited and Quoted 0%
Matches with in-text citation present, but no quotation marks

Top Sources

7% Internet sources
4% Publications
9% Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

Appendix B: Ethical Clearance Confirmation

Strathmore UNIVERSITY

1st October 2024

Mr Kiundi Richard,
richard.kiundi@strathmore.edu

Dear Mr Kiundi,

RE: Dynamic Denial of Service Attack Prevention using a Multi-Level IP Shedding Defense Mechanism

This is to inform you that SU-ISERC has reviewed and **approved** your above **SU-masters** proposal. Your application reference number is **SU-ISERC2391/24.** The approval period is from **1st October 2024 to 30th September 2025.**

This approval is subject to compliance with the following requirements:

i. Only approved documents including (informed consents, study instruments, MTA) will be used.
ii. All changes including (amendments, deviations, and violations) are submitted for review and approval by SU-ISERC.
iii. Death and life-threatening problems and serious adverse events or unexpected adverse events whether related or unrelated to the study must be reported to SU-ISERC within 72 hours of notification.
iv. Any changes anticipated or otherwise that may increase the risks or affected safety or welfare of study participants and others or affect the integrity of the research must be reported to SU-ISERC within 72 hours.
v. Clearance for the export of biological specimens must be obtained from relevant institutions.
vi. Submission of a request for renewal of approval at least 60 days prior to the expiry of the approval period. Attach a comprehensive progress report to support the renewal.
vii. Submission of an executive summary report within 90 days of completion of the study to SU-ISERC.

Before commencing your study, you will be expected to obtain a research license from National Commission for Science, Technology, and Innovation (NACOSTI) https://research-portal.nacosti.go.ke/ and obtain other clearances needed.

Yours sincerely,

Mr Ambrose Rachier,
Chairperson; SU-ISERC

Ole Sangale Rd, Madaraka Estate. PO Box 59857-00200, Nairobi, Kenya. Tel +254 (0)703 034000
Email admissions@strathmore.edu www.strathmore.edu

Appendix C: PHP code for IP Location Classifier Component

```
<?php
require("connection.php");

function getIPLocation($ip) {
   //$url = "http://ip-
api.com/json/{$ip}?fields=status,message,country,regionName,city,lat,lon,isp,
org,as";

   $url = "http://ip-
api.com/json/{$ip}?fields=status,message,country,regionName,city,isp,org,as"
;
   //$url = "http://ip-api.com/json/{$ip}?fields=status,message,country,as";
   $response = file_get_contents($url);

   if ($response) {
      $data = json_decode($response, true);
      if ($data['status'] === 'success') {
         return $data;
      } else {
         return "Error: " . $data['message'];
      }
   } else {
      return "Error fetching data.";
   }
}

// Function to determine RIR from AS number
function get_rir($as) {
   $as_number = intval(preg_replace('/[^0-9]/', '', $as)); // Extract numeric AS
number

   if ($as_number >= 1 && $as_number <= 399999) {
      return "ARIN (North America)";
   } elseif ($as_number >= 400000 && $as_number <= 599999) {
      return "RIPE NCC (Europe, Middle East, Central Asia)";
   } elseif ($as_number >= 600000 && $as_number <= 799999) {
      return "APNIC (Asia-Pacific)";
   } elseif ($as_number >= 800000 && $as_number <= 999999) {
      return "LACNIC (Latin America, Caribbean)";
   } elseif ($as_number >= 1000000 && $as_number <= 1999999) {
      return "AFRINIC (Africa)";
   } else {
      return "Unknown RIR";
   }
}

function getZone($country){
   $zone = 2;
   if(strcmp(strtolower(trim($country)), "kenya") == 0){
```

```
        $zone = 1;
    }

    return $zone;
}

// Example usage
//$ip0 = "8.8.8.8"; // Replace with any public IPv4 address
$ip1 = "10.20.113.39";
$ip2 = "34.243.183.166";
$ip3 = "156.0.233.51";
$ip4 = "1.1.1.1";

$ip = $ip4;

$location = getIPLocation($ip);
//print_r($location);

if (is_array($location)) {
    $country = secureString($location['country']);
    $zone = getZone($country);
    //$region = secureString($location['regionName']);
    $city = secureString($location['city']);
    //$latitude = secureString($location['lat']);
    //$longitude = secureString($location['lon']);
    $isp = secureString($location['isp']);
    $organization = secureString($location['org']);
    $as = secureString($location['as']);
    $rir = get_rir($as);

    // Output values
    echo "IP: $ip\n";
    echo "Country: $country\n";
    echo "Zone: $zone\n";
    //echo "Region: $region\n";
    echo "City: $city\n";
    //echo "Latitude: $latitude\n";
    //echo "Longitude: $longitude\n";
    echo "ISP: $isp\n";
    echo "Organization: $organization\n";
    echo "AS: $as\n";
    echo "RIR: $rir\n";

    $Connection = mysqli_connect($Host, $User, $Password, $dummyDB) or
die ("Error 0 - Unable to connect. Check your connection parameters.");
    $query = "INSERT INTO `ipZones`.`zonedata`
(`ipAddress`,`zone`,`country`,`city`,`asn`,`ISP`,`LIR`,`RIR`)
    VALUES ('$ip','$zone','$country','$city','$as','$isp','$organization','$rir')";

    mysqli_query($Connection, $query) or die(mysqli_error($Connection));
```

```
    mysqli_close($Connection);

} else {
    echo $location; // Display error message if API call fails
}

?>
```

Appendix D: PHP code for Profit Classifier Component

```
<?php
require("connection.php");

function getIPLocation($ip) {
    $url = "http://ip-api.com/json/{$ip}?fields=status,message,country";

    $response = file_get_contents($url);

    if ($response) {
        $data = json_decode($response, true);
        if ($data['status'] === 'success') {
            return $data;
        } else {
            return "Error: " . $data['message'];
        }
    } else {
        return "Error fetching data.";
    }
}

function getZone($country){
    $zone = 2;
    if(strcmp(strtolower(trim($country)), "kenya") == 0){
        $zone = 1;
    }

    return $zone;
}

$ip2 = "34.243.183.166";
$ip3 = "156.0.233.51";

$ip = $ip3;
$value = 100;
$time = date("Y-m-d H:i:s");
```

```
$location = getIPLocation($ip);

if (is_array($location)) {
    $country = secureString($location['country']);
    $zone = getZone($country);

    // Output values
    /*
    echo "IP: $ip\n";
    echo "Country: $country\n";
    echo "Zone: $zone\n";
    */

    //Log zone requests
    $Connection = mysqli_connect($Host, $User, $Password, $dummyDB) or
die ("Error 0 - Unable to connect. Check your connection parameters.");
    $query = "INSERT INTO `ipZones`.`revenueData` (`zone`,`value`,`time`)
    VALUES ('$zone','$value','$time')";
    mysqli_query($Connection, $query) or die(mysqli_error($Connection));
    mysqli_close($Connection);

    //Update zone profits
    $Connection = mysqli_connect($Host, $User, $Password, $dummyDB) or
die ("Error 0 - Unable to connect. Check your connection parameters.");
    $query = "UPDATE `ipZones`.`profitabilityData` SET
`value`=`value`+$value
    WHERE `zone`='$zone'";
    mysqli_query($Connection, $query) or die(mysqli_error($Connection));
    mysqli_close($Connection);

} else {
    echo $location; // Display error message if API call fails
}
```

Appendix E: Bash Script To Measure Server Load

```
#!/bin/bash

# Function to get the total number of connections to Apache
get_total_connections() {
   ss -ant | grep ':80\|:443' | grep ESTAB | wc -l
}

# Function to get the number of requests currently being processed by Apache
get_active_requests() {
   apachectl status 2>/dev/null | grep 'BusyWorkers' | awk '{print $2}'
}

# Function to get the number of requests currently being processed by Apache
get_active_tasks() {
   systemctl status apache2 | grep 'Tasks' | awk '{print $2}'
}

# create a loop running for 1 minute
Num=0
while [ $Num -lt 300 ]; do
   TOTAL_CONNECTIONS=$(get_total_connections)
   #ACTIVE_REQUESTS=$(get_active_requests)
   ACTIVE_TASKS=$(get_active_tasks)

   #echo "Total active connections: $TOTAL_CONNECTIONS"
   #echo "Total active requests (Tasks): $ACTIVE_TASKS"

   echo "Connections: $TOTAL_CONNECTIONS, Tasks:
$ACTIVE_TASKS"
   echo "$Num, $TOTAL_CONNECTIONS, $ACTIVE_TASKS" >>
data1.csv

   Num=$((Num + 1))  # Correct way to increment the loop variable
   sleep 1  # Ensures the loop runs every second
done
```

Appendix F: PHP Script for Client Request Tool

```
<?php

// Web server URL (replace with your actual URL)
$server_url = "http://172.16.13.10/app.php";

// Generate a random number between 10 and 5000
$value = rand(5, 100);

// Build the request URL with query parameters
$request_url = $server_url . "?value=" . urlencode($value);

// Initialize cURL session
$ch = curl_init();

// Set cURL options
curl_setopt($ch, CURLOPT_URL, $request_url);
curl_setopt($ch, CURLOPT_RETURNTRANSFER, true);
curl_setopt($ch, CURLOPT_TIMEOUT, 2); // Set timeout to 5 seconds
// Ensure a new connection is used for each request
curl_setopt($ch, CURLOPT_FRESH_CONNECT, true);
curl_setopt($ch, CURLOPT_FORBID_REUSE, true);

// Add HTTP header to explicitly close the connection
curl_setopt($ch, CURLOPT_HTTPHEADER, ["Connection: close"]);

// Execute the request
$response = curl_exec($ch);

// Check for errors
if ($response === false) {
   /*
   $errorMessage = date("Y-m-d H:i:s") . " - Request timed out or failed for
value: $value - Error: " . curl_error($ch) . "\n";
   echo "Request failed. Logged to request_log.txt\n";
   */
   $errorMessage = date("Y-m-d H:i:s") . ", 1\n";

   // Log the error to a file
   file_put_contents("clientFail1.csv", $errorMessage, FILE_APPEND);
} else {
   /*
   echo "Sent value: $value\n";
```

```
    echo "Server response: $response\n";
    */

    $respMessage = date("Y-m-d H:i:s") . ", 1\n";
    // Log the error to a file
    file_put_contents("clientResp1.csv", $respMessage, FILE_APPEND);
}

// Close cURL session
curl_close($ch);

?>
```

Appendix G: Web server application code

```
<?php
require("connection.php");

$ip = $_SERVER['REMOTE_ADDR'];

if(isset($_GET["value"])){
    $value = secureString($_GET['value']);
} else {
    $value = 0;
}

if(!is_numeric($value)){
    $value = 0;
}

function getZoneVal($ip){
    $zoneArray = [
        "172.16.13.20"=>"1",
        "172.16.13.62"=>"1",
        "172.16.13.30"=>"2"
    ];

    $zone = 2;
    $zone = $zoneArray[$ip];

    return $zone;
}

$zone = getZoneVal($ip);
$time = date("Y-m-d H:i:s");
```

```
//Log zone requests
$Connection = mysqli_connect($Host, $User, $Password, $dummyDB) or die
("Error 0 - Unable to connect. Check your connection parameters.");
$query = "INSERT INTO `ipZones`.`revenueData` (`zone`,`value`,`time`)
VALUES ('$zone','$value','$time')";
mysqli_query($Connection, $query) or die(mysqli_error($Connection));
mysqli_close($Connection);

//Update zone profits
$Connection = mysqli_connect($Host, $User, $Password, $dummyDB) or die
("Error 0 - Unable to connect. Check your connection parameters.");
$query = "UPDATE `ipZones`.`profitabilityData` SET `value`=`value`+$value
WHERE `zone`='$zone'";
mysqli_query($Connection, $query) or die(mysqli_error($Connection));
mysqli_close($Connection);

sleep(1);
echo "200 ok";
?>
```

Appendix H: Bash script to activate iptables rules

```
#!/bin/bash

# IP address to block
BLOCK_IP="172.16.13.30"

# Flush existing rules for this IP to avoid duplicates
sudo iptables -D INPUT -s $BLOCK_IP -j DROP 2>/dev/null

# Add iptables rule to drop packets from the specified IP
# Adding to the beginning of the chain with higher priority
sudo iptables -I INPUT 1 -s $BLOCK_IP -j DROP

# Save the rules to make them persistent after reboot
if command -v iptables-save >/dev/null 2>&1; then
   case "$(lsb_release -is 2>/dev/null)" in
      "Debian"|"Ubuntu")
         sudo iptables-save | sudo tee /etc/iptables/rules.v4 >/dev/null
         ;;
      "CentOS"|"RedHat"|"Fedora")
         sudo service iptables save
         ;;
```

```
        *)
            echo "Warning: Rules may not persist after reboot. Manual save
required."
            ;;
    esac
fi

# Verify the rule was added
echo "Blocked all incoming packets from $BLOCK_IP"
echo "Current iptables rules:"
sudo iptables -L INPUT -n --line-numbers | grep $BLOCK_IP
```