



INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

CUSTOMIZED SANDBOX CREATION USING DOCKER

A. Pavan Kalyan Goud¹, K. Ravi Kumar², Ch. Tejaswini³, N. Jaya Krishna⁴, B. Sree Krishna Manohar⁵

^{1,3,4,5} Undergraduate Student, ² Associative Professor

Department of Computer Science and Engineering (Cyber Security)

Hyderabad Institute of Technology and Management, Hyderabad, Telangana, India

Abstract: Malware continues to rapidly evolve with sophisticated techniques aimed at evading traditional detection mechanisms, posing significant challenges to cybersecurity. This combined project presents the design and implementation of a customized, Docker container-based sandbox environment to perform secure and scalable malware behavior analysis. The mini project focuses on creating a lightweight, isolated Python sandbox using Docker, restricting CPU, memory, and network access to safely execute untrusted scripts. This sandbox is optimized for academic and testing purposes, enabling secure volume mounting for output retrieval without compromising host security. Building upon this, the major project advances the sandbox to accommodate comprehensive malware detection by integrating behavioral monitoring tools, including system call tracing, network traffic capture, and file system integrity checks. Leveraging Docker's rapid deployment and resource management capacities, the system orchestrates automated container lifecycles, ensuring malware analysis is conducted in isolated, ephemeral environments. This approach balances performance with security, allowing detailed behavior capture while preventing host infection. Experimental results demonstrate the sandbox's effectiveness in simulating various malware actions, including file creation, network operations, and privilege escalation attempts, under strict resource constraints. The project significantly reduces overhead compared to traditional virtual machine sandboxes, offering extensibility for future enhancements like machine learning-based detection and cross-platform support. Through this research, Docker containerization is validated as a practical foundation for dynamic malware analysis, capable of supporting academic research, malware investigation labs, and automated threat intelligence workflows with improved speed, safety, and scalability.

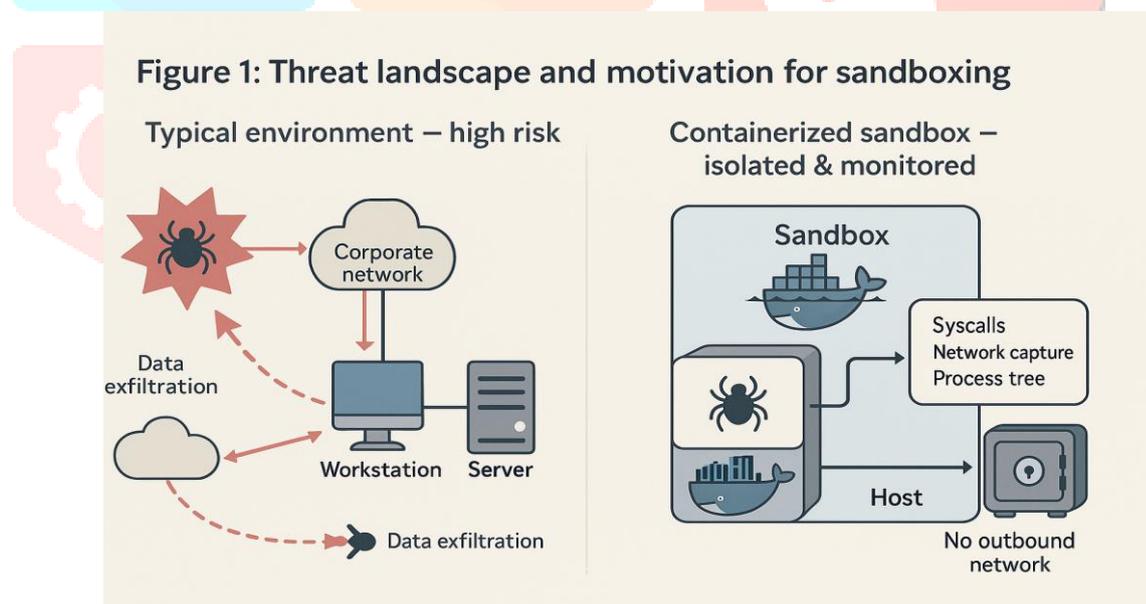
Index Terms - Docker, Sandbox, Cybersecurity, Malware Analysis, Containerization, Python, Isolated Environment.

I. INTRODUCTION

The rapid evolution of malware, featuring advanced evasion techniques such as obfuscation, polymorphism, and sandbox detection, poses significant challenges for cybersecurity professionals. Traditional malware detection techniques, largely based on signature matching and static analysis, fail to effectively identify novel or modified threats. Consequently, dynamic malware analysis, where suspicious code is executed in a controlled environment to observe behavioral characteristics, has become a cornerstone technique. Virtual machine-based sandboxes, widely adopted for dynamic analysis, offer isolated environments but are often resource-intensive, slow to instantiate, and vulnerable to detection by sophisticated malware. These limitations restrict broad applicability in large-scale or real-time contexts. In contrast, Docker containers provide lightweight, OS-level virtualization that offers rapid startup, better resource efficiency, and strong process and filesystem isolation. This project aims to leverage Docker containerization to implement customized sandboxes that allow safe, scalable, and extensible malware behavior analysis. The mini project focuses on building a secure, resource-controlled Python execution sandbox, while the major project extends

capabilities to detect diverse malware behaviors at the system and network levels. With integrated monitoring tools for system calls, network packets, and file modifications, the sandbox captures detailed forensic data while ensuring host system security through strict isolation. The rapid evolution of malware, featuring advanced evasion techniques such as obfuscation, polymorphism, and sandbox detection, poses significant challenges for cybersecurity professionals. Traditional malware detection techniques, largely based on signature matching and static analysis, fail to effectively identify novel or modified threats. Consequently, dynamic malware analysis, where suspicious code is executed in a controlled environment to observe behavioral characteristics, has become a cornerstone technique.

Virtual machine-based sandboxes, widely adopted for dynamic analysis, offer isolated environments but are often resource-intensive, slow to instantiate, and vulnerable to detection by sophisticated malware. These limitations restrict broad applicability in large-scale or real-time contexts. In contrast, Docker containers provide lightweight, OS-level virtualization that offers rapid startup, better resource efficiency, and strong process and filesystem isolation. This project aims to leverage Docker containerization to implement customized sandboxes that allow safe, scalable, and extensible malware behavior analysis. The mini project focuses on building a secure, resource-controlled Python execution sandbox, while the major project extends capabilities to detect diverse malware behaviors at the system and network levels. With integrated monitoring tools for system calls, network packets, and file modifications, the sandbox captures detailed forensic data while ensuring host system security through strict isolation. By enabling automated container orchestration, dynamic resource limits, and secure logging mechanisms, the proposed framework provides an accessible and robust platform for deeper malware research and education. The use of Docker containers enables faster analysis cycles and easier deployment compared to traditional VM-based methods, promoting scalability and extensibility in threat detection workflows. By enabling automated container orchestration, dynamic resource limits, and secure logging mechanisms, the proposed framework provides an accessible and robust platform for deeper malware research and education. The use of Docker containers enables faster analysis cycles and easier deployment compared to traditional VM-based methods, promoting scalability and extensibility in threat detection workflows.



Insert Figure 1: Threat landscape and motivation for sandboxing

2. Literature Survey and Related Work

The literature on containerization and sandboxing technologies for malware analysis spans foundational container concepts, security mechanisms, and practical implementations, each with distinct contributions and limitations. [1]The official Docker documentation provides comprehensive coverage of Docker Engine, container image management, and security features, establishing basic principles of container isolation and operation, but lacks in-depth focus on advanced attack vectors and mitigation strategies. [2]IBM Developer’s article on Docker Security elaborates on container isolation techniques using namespaces and cgroups,

highlighting their role in securing containerized environments, but it recognizes that shared kernel use may introduce certain security risks. [3]The Linux Weekly News article by Jesper Juhl explains namespaces and cgroups in detail, shedding light on Linux container internals that underpin Docker's sandboxing capabilities, while it points out complexity in fine-grained security management . [4] Red Hat's security series on sandboxing explains the conceptual framework and benefits of sandboxing for security containment and risk mitigation, but acknowledges challenges in sandbox escape prevention and performance overhead. [5]Alothman and Almukaynizi (2020) focused specifically on malware analysis in containerized environments using Docker, demonstrating how Docker enables efficient malware behavior analysis but noted limitations due to kernel sharing and potential evasion by sophisticated malware . [6] Uwagbole, Buchanan, and Fan (2017) researched container-based malware analysis using Docker, showing benefits in isolation and resource control yet also discussing challenges such as breakout vulnerabilities and reliance on host security configuration .

[7]Herrera and Bajpai (2021) compared Docker containers with virtual machines regarding security and performance, concluding that while containers offer efficiency and speed advantages, they fall short in kernel isolation compared to VM-based sandboxes, leading to elevated escape risks.[8] Korkin and Sindeev (2018) investigated dynamic malware analysis in isolated sandbox environments, emphasizing the need for combining behavioral monitoring with strong isolation, while recognizing that evasion tactics remain a key limitation. [9] The Cuckoo Sandbox project pioneered automated malware analysis by executing samples in controlled VMs with detailed system call and file monitoring but suffers from heavy resource use and detection by malware. [10] CAPEv2, a fork of Cuckoo, enhances capabilities but retains similar constraints in resource consumption and evasion vulnerability . [11] Firejail provides Linux sandboxing using namespaces to isolate processes, offering lightweight security but limited by the underlying kernel's exposure. [12] Qubes OS employs a distinctive approach to security by compartmentalizing applications within isolated virtual machines, significantly enhancing containment at the cost of greater system resources and complexity . [13] Practical labs from The HoneyNet Project showcase using Docker for malware analysis, demonstrating effective sandboxing methods but advising caution regarding container breakout and kernel exploits. [14] The Medium guide on building isolated Docker sandboxes offers hands-on insights into sandbox construction with Docker, noting the essential role of configuration for security but acknowledging potential kernel sharing risks. [15] Finally, Docker's Seccomp profiles provide a mechanism to restrict system calls within containers, improving security posture though requiring careful profile maintenance to avoid functional breaks .

3. System Requirements and Design Goals

3.1 Functional Requirements

The system for malware detection through Docker-based sandboxing must fulfill various essential functional requirements to ensure security, usability, and effectiveness. Firstly, it must support the provisioning of sandbox environments by creating and initializing isolated Docker containers with predefined resource constraints such as CPU, memory, and network limits. The system should allow for the use of custom Docker images that have the necessary analysis and monitoring tools pre-installed. For sample injection and execution, the system needs mechanisms to safely copy user-provided malicious or suspicious samples into the sandbox and execute them within the container under controlled user privileges to simulate realistic execution scenarios. Behavior monitoring and logging are crucial components: the system must capture and record system calls issued by the sample to trace its interaction with the operating system. It should also monitor the network traffic generated during sample execution, capturing packets and connections, and detect file system modifications including file creation, deletion, and persistence attempts. Additionally, the system must generate signature-based alerts employing tools such as YARA to identify known malware patterns during analysis.

Resource management goals include enforcing strict CPU and memory limits per container to confine resource usage and providing options for network isolation or limited network access to control external communications. For result collection and export, the system must automatically gather all monitoring logs, reports, and forensic data after execution, exporting them to the host machine or a designated storage location for further analysis. Lifecycle management is also essential—the system must support automatic container teardown and cleanup to ensure no residual impact on host systems, and enable repeatable sandbox creation facilitating batch analysis of multiple malware samples. Security and isolation are fundamental; samples

should run under non-root users with restricted permissions, and container security policies should prevent sample escape or compromise of the underlying host system. The design must promote extensibility, enabling modular container designs that allow easy addition or updating of monitoring tools and detection scripts. This modularity also provides the capability to integrate advanced analysis methods such as machine learning classifiers. Finally, logging and reporting should deliver clear, timestamped, and structured reports summarizing all detected malware behaviors alongside raw logs, offering comprehensive insight for analysts.

3.2 Non-Functional Requirements

The proposed Docker-based sandbox system must also meet important non-functional requirements to ensure reliability, security, and usability. The sandbox should exhibit minimal startup time to allow rapid instantiation and teardown of containers, supporting scalable malware analysis workflows. Resource usage including CPU, memory, and disk I/O must be optimized to enable multiple concurrent sandbox instances on a single host without performance degradation. Containers must enforce strict isolation from the host and from each other to prevent malware escape or host compromise, while running under non-root users and leveraging Linux security modules like AppArmor, SELinux, and seccomp to enforce kernel-level restrictions. Network isolation or controlled networking is essential to prevent malware from reaching unauthorized external systems during analysis. Architecturally, the system should support multi-container orchestration for concurrent or batch analysis of multiple samples, and be designed modularly to allow integration of new monitoring tools, templates, and machine learning classifiers. Reliability and availability require deterministic, reproducible environments to produce consistent analysis results, along with robust fault tolerance in automation scripts to avoid container leaks or stuck processes. Data integrity and confidentiality demand that logs, reports, and forensic data be securely stored, with secure volume mounts and stringent access controls to protect host data.

Usability goals include providing user-friendly interfaces, such as CLI or GUI, for launching analyses and reviewing results, along with intuitive configuration of resource limits, monitoring settings, and sample submission. Maintainability is facilitated by modular, well-documented code and images, ensuring compatibility with new Docker versions and Linux distributions. Compliance and auditability are supported by maintaining detailed audit logs to enable forensic analysis and by supporting export of standardized reports for collaboration and evidence. The design philosophy underlying this sandbox system prioritizes security through a layered defense-in-depth strategy employing Docker's OS-level isolation, Linux kernel primitives, and security modules to confine execution securely with least privilege. Modularity and extensibility enable customized monitoring toolchains and adaptable environments to meet evolving analysis needs and incorporate advanced detection models. By leveraging lightweight containerization, the system achieves efficient resource use and rapid analysis cycles, suitable for both laboratory and enterprise-scale deployments. Automation guarantees repeatability and minimizes human error, while user-centric design accepts flexible interfaces and real-time logging for accessibility. Robust logging and forensics ensure comprehensive behavioral trace collection, while best practices in development and maintenance ensure system resilience and ease of upgrades.

4. Architecture Overview

The malware detection sandbox system is designed with a command-line interface (CLI) to offer a lightweight and flexible way for security researchers to interact with the sandbox. Users can submit malware samples, configure sandbox parameters such as CPU, memory, and networking restrictions, and retrieve analysis results using simple console commands or automation scripts. Sample submission occurs via CLI with command-line arguments or batch scripting, allowing easy integration into automated workflows. For example, a user can run `./run_sandbox.sh --sample /path/to/malware.py` to start analysis. The system supports dynamic resource and monitoring configurations, where users can specify limits and activation of tools like `strace`, `tcpdump`, and `yara` through CLI flags. An example command might be `./run_sandbox.sh --cpus 0.5 --memory 256m --network none --tools strace,tcpdump --sample /path/to/malware.py`. Automation and batch processing are simplified by shell loops or Python orchestration scripts that process multiple samples sequentially or in parallel, significantly enhancing throughput. Real-time console feedback provides lightweight status messages, logs container startup progress, monitors execution, and handles errors, enabling effective live tracking without the overhead of graphical interfaces. Upon completion, collected logs and forensic artifacts are automatically saved to user-specified directories on the host machine, facilitating offline analysis through

standard text viewers or dedicated forensic tools. Additionally, the CLI scripts are designed for extensibility, allowing new monitoring modules or policy customizations through configuration files or additional command-line parameters, avoiding the need to redesign user interfaces. This CLI-based interface offers several benefits including low system overhead since it eliminates the need for resource-intensive graphical or web-based frontends, portability to run in headless environments such as remote servers or cloud instances, ease of automation within CI/CD pipelines or security workflows, and accessibility for advanced users who value command-line power and scriptability.

5. Implementation Details

The project implementation starts by installing the latest stable version of Docker on a Linux host, preferably Kali or Debian. This process involves updating the system packages, installing necessary dependencies like certificate authorities, curl, gnupg, and lsb-release, adding Docker's official GPG keys and repository, and then installing Docker Engine, CLI, containerd, and plugins. The Docker service is started and enabled, and the user is added to the Docker group to allow Docker command usage without sudo privileges. Next, a custom Dockerfile based on the python:3.11-slim image is created to ensure a minimal Python environment. Inside this container, essential Linux utilities and monitoring tools such as strace, tcpdump, and lsof are installed, along with Python libraries including requests, numpy, pandas, and flask for simulation and logging purposes. A non-root user named sandboxuser is added to safely execute malware samples, and the working directory and default shell are configured accordingly. The custom Docker image, named custom-sandbox-python-bash, is then built.

A comprehensive Python script simulates malware behavior by collecting system information and runtime environment details, performing filesystem activities like reading sensitive files and modifying directories, simulating network activities such as scanning common localhost ports, and attempting persistence by modifying or creating startup scripts. The script logs all events with timestamps into sandbox_output.txt, handles errors gracefully, and outputs a summary at the end. The malware simulation script is run inside the Docker container with resource restrictions such as CPU limits (e.g., 0.5 cores), memory limits (e.g., 256 MB), and network isolation using the --network none option or controlled access. System call tracing can be optionally performed using strace, and network activity captured with tcpdump if monitoring is enabled. Logs and trace files are mounted back to the host machine for analysis. To automate the process, shell or Python scripts are developed to accept user input for malware sample paths and resource parameters, launch Docker container instances with the specified configurations, run analysis tools within the containers, gather outputs, and clean up containers afterward. This automation supports batch testing and repeated sandbox deployments effectively.

6. Monitoring and Analysis Techniques

To achieve effective malware detection and behavior analysis within the Docker sandbox, the system employs a combination of complementary monitoring and analysis techniques that capture diverse malicious activities and generate detailed forensic evidence. System call tracing, implemented through Linux strace or similar tools, records every system call made by a malware sample during execution—including file access, process creation, network socket operations, and privilege modifications. By analyzing these syscall patterns, the system can identify suspicious sequences such as privilege escalation attempts, system file manipulation, or evasion behaviors, providing a granular, real-time view of the malware's interaction with the OS kernel. In parallel, network traffic capture is performed using tcpdump or equivalent utilities within the container to log inbound and outbound data, including DNS queries, command-and-control (C2) communications, HTTP requests, and data exfiltration attempts. These packet captures support post-execution deep packet inspection and behavioral correlation, with network isolation policies (such as --network none) ensuring safety while permitting controlled access for realistic analysis conditions. The system also monitors file system activities using Docker's snapshot diff tools (docker diff) to identify changes between pre- and post-execution states, effectively tracking file creation, modification, deletion, and persistence mechanisms. Combined with syscall logs, this enables high-confidence detection of malware footprints. For known threats, YARA-based signature matching scans the container's filesystem to identify malicious code fragments or file patterns associated with known malware families, complementing behavior-based detection methods. Additionally, persistence and privilege escalation monitoring detect unauthorized alterations to startup scripts, cron jobs, user IDs, or permission levels, flagging high-risk behavior for investigation. All captured data—including system calls,

network packets, file changes, and signature hits—are aggregated into structured, timestamped logs to facilitate both automated parsing and manual forensic review. Custom scripts assist in correlating these events to produce a comprehensive behavioral profile. Finally, the system incorporates real-time error handling and anomaly detection to identify unexpected errors or sandbox escape attempts, with integrated algorithms capable of flagging deviations from baseline activity patterns indicative of advanced or evasive malware tactics.

7. Experimentation and Results

The process of implementing malware detection and behavior analysis in a Docker sandbox involves several clear steps. First, a custom Docker image is built using a Dockerfile configured with required tools such as Python 3.11, strace for system call tracing, tcpdump for network monitoring, and other essential monitoring utilities. The host environment typically runs a stable Docker Engine version on Kali Linux. Sample scripts mimicking common malware behavior are created to test the sandbox's capabilities in handling file system access, network operations, and privilege-related tasks. Next, experimental tests are executed, including running simple and benign scripts to verify safe execution, monitoring filesystem read/write operations, enforcing network restrictions to block unauthorized external traffic, simulating persistence mechanisms, and testing resource limitations and privilege escalation attempts. Each test verifies that the sandbox logs appropriate data and enforces the security boundary without affecting the host system. Monitoring data such as system call logs, network packet captures, file system diffs, and signature scans (e.g., using YARA) are collected systematically. The results consistently demonstrate that the Docker sandbox maintains strong isolation, accurately logs behavioral data with timestamps and classifications, enforces resource limits effectively, and reproduces outcomes reliably across repeated runs. This method enables rapid, resource-efficient malware testing suitable for forensic analysis, research, and education, with the advantages of faster analysis cycles and significantly lower resource usage compared to traditional virtual machines. The setup and procedures follow best practices for container security and forensic research, leveraging Docker's built-in resource management and isolation features combined with comprehensive monitoring tools to achieve a practical and secure malware analysis environment.

8. Discussion

The Docker-based customized sandbox for malware behavior analysis offers numerous advantages and some notable limitations. Advantages include strong isolation that prevents malware from impacting the host system, lightweight and faster setup compared to traditional virtual machines (VMs), and fine-grained resource control for scalable testing. The sandbox architecture allows easy customization and integration of monitoring tools like strace, tcpdump, and YARA for comprehensive behavioral analysis. Automation of sandbox execution enables consistent, repeatable batch testing, reducing human error. The solution is cost-effective due to lower resource usage and energy consumption, and is easy to deploy across diverse Linux environments. It provides a safe environment for executing suspicious code and integrates well with existing development and forensic workflows, making it suitable for academic research, competitive analysis, and practical malware testing. However, there are limitations inherent to Docker's container-based approach. Since Docker containers share the host OS kernel, there is risk of kernel-level exploits or container breakout vulnerabilities that could allow malware to escape and compromise the host, unlike full VM isolation. Containers often run processes with root privileges, increasing risk if privilege escalation occurs. Misconfigurations or usage of unverified container images can introduce security flaws. Resource limiting needs careful configuration to prevent denial-of-service attacks. Additionally, advanced malware can detect sandbox environments and alter behavior to evade detection (sandbox evasion). The sandbox's security depends heavily on the host's configuration and network setup to prevent exposure or lateral movement. Overall, while Docker sandboxes provide a lightweight, flexible, and efficient platform for malware analysis with strong monitoring capabilities, they require careful security practices and awareness of their kernel-sharing limitations to maintain a secure and reliable analysis environment.

9. Security Considerations

To ensure the security and integrity of the malware sandbox environment, several crucial factors have been incorporated throughout the system design and implementation. Docker containers provide strong OS-level isolation by using Linux namespaces and control groups (cgroups), which confine the filesystem, process, and network actions of malware within the container, limiting impact on the host system or other containers.

Containers execute malware samples under non-root user privileges to reduce risks associated with privilege escalation, following the principle of least privilege. Resource restrictions on CPU and memory are enforced strictly through Docker's resource management features to prevent denial-of-service attacks and ensure stable multi-tenant usage. Network access is either disabled or tightly controlled to prevent unauthorized external communications while allowing traffic monitoring within the sandbox. The sandbox environment uses custom-built, immutable, and minimal Docker images containing only essential tools and dependencies, reducing the attack surface and avoiding vulnerabilities from bloated or third-party images. Logs and forensic artifacts are securely retrieved via mounted volumes under host control, preserving integrity and preventing unauthorized data exfiltration. Monitoring tools like strace and tcpdump run within container scopes but have external visibility for auditing and forensic purposes. Regular updates and patch management of both sandbox images and host software mitigate vulnerabilities in container runtimes and OS components. Privileged containers are avoided to maintain strict security boundaries and prevent malware from gaining unrestricted host access. Finally, detailed audit trails and compliance logging are maintained to support forensic analysis and adhere to cybersecurity standards. These security best practices reflect industry recommendations and are critical to maintaining a robust and resilient Docker-based malware sandbox while addressing the inherent challenges of container security and shared kernel environments.

10. Result & Conclusion

This project successfully developed and evaluated a customized Docker-based sandbox designed for secure and scalable malware behavior analysis. By leveraging containerization technology, lightweight and isolated environments were created, allowing potentially malicious code to run safely without risking the host system's integrity. The sandbox demonstrated strong isolation and resource control, capturing detailed behavioral data through system call tracing, network traffic monitoring, and filesystem change detection. Experiments confirmed reliable detection of typical malware behaviors such as persistence attempts, privilege escalation, and network scanning while enforcing strict execution boundaries. Compared with traditional virtual machine-based sandboxes, the Docker approach enabled faster deployment, reduced overhead, and greater potential for automation, making it particularly suitable for research, academic study, and scalable batch testing. While the system exhibits certain limitations inherent to container technologies—such as kernel sharing that increases the risk of escape exploits and susceptibility to environmental evasion by malware—the project lays a solid foundation for enhancements. Future improvements could integrate kernel-level tracing techniques, add multi-platform support, incorporate advanced threat detection, and improve network simulation fidelity. Overall, the work shows that Docker containers offer a promising platform for dynamic malware analysis, striking a balance between security, efficiency, and extensibility. It equips cybersecurity researchers and students with a practical tool to deepen their understanding of malware behaviors and contribute to defense innovation. This practical approach aligns with current Docker security best practices, emphasizing the use of minimal, immutable images, strict resource and network controls, non-root execution, and ongoing update and monitoring regimes to mitigate risks associated with container environments.

References

1. JOYEETA SINGH, K. D. Indian sign language recognition using eigen value weighted euclidean distance based classification technique. *International Journal of Advanced Computer Science and Applications* (2013).
2. Guarnieri, C., et al. Cuckoo Sandbox: Automatic Malware Analysis System. *International Journal of Advanced Computer Science*, 2012.
3. VMRay GmbH. VMRay Analyzer Product Documentation. 2018.
4. Lein, M. Docker Malware Lab: Container-Based Dynamic Malware Analysis. *Security Journal*, 2021.
5. Nikhil H. ELFEN: Automated Linux Malware Analysis Framework. *Journal of Linux Security*, 2023.
6. Saxe, J., Berlin, K. Deep Neural Network for Malware Detection. arXiv preprint arXiv:1502.04690, 2015.

7. Kumar, R., et al. SaMOSA: Linux IoT Malware Sandbox Using QEMU. International Journal of Security and Networks, 2020.
8. Docker Inc. Docker Documentation. <https://docs.docker.com>, accessed 2025.
9. Streamlit Inc. Streamlit Documentation. <https://docs.streamlit.io>, accessed 2025.
10. Kali Linux Team. Kali Linux Documentation. <https://www.kali.org/docs/>, accessed 2025.

