

HoneyWin: High-Interaction Windows Honeypot in Enterprise Environment

Yan Lin Aung*
y.aung@derby.ac.uk
University of Derby
Derby, United Kingdom

Yee Loon Khoo
Davis Yang Zheng[†]
Bryan Swee Duo[†]
Singapore Institute of Technology
Singapore

Sudipta Chattopadhyay
Jianying Zhou
Singapore University of Technology and Design
Singapore

Liming Lu
Weihan Goh
Singapore Institute of Technology
Singapore

ABSTRACT

Windows operating systems (OS) are ubiquitous in enterprise Information Technology (IT) and operational technology (OT) environments. Due to their widespread adoption and known vulnerabilities, they are often the primary targets of malware and ransomware attacks. With 93% of the ransomware targeting Windows-based systems, there is an urgent need for advanced defensive mechanisms to detect, analyze, and mitigate threats effectively. In this paper, we propose **HoneyWin** a high-interaction Windows honeypot that mimics an enterprise IT environment. The **HoneyWin** consists of three Windows 11 endpoints and an enterprise-grade gateway provisioned with comprehensive network traffic capturing, host-based logging, deceptive tokens, endpoint security and real-time alerts capabilities. The **HoneyWin** has been deployed live in the wild for 34 days and receives more than 5.79 million unsolicited connections, 1.24 million login attempts, 5 and 354 successful logins via remote desktop protocol (RDP) and secure shell (SSH) respectively. The adversary interacted with the deceptive token in one of the RDP sessions and exploited the public-facing endpoint to initiate the Simple Mail Transfer Protocol (SMTP) brute-force bot attack via SSH sessions. The adversary successfully harvested 1,250 SMTP credentials after attempting 151,179 credentials during the attack.

CCS CONCEPTS

• **Security and privacy** → **Intrusion detection systems; Malware and its mitigation.**

KEYWORDS

High-Interaction Windows Honeypot, Deception, Network Traffic Analysis, Host Log Analysis, Attack Attribution

1 INTRODUCTION

Windows operating systems (OS) are extensively utilized in enterprise Information Technology (IT) and operational technology (OT) environments, while their prevalence and roles differ across these domains. In 2025, Windows OS maintains a dominant position holding more than 70% of the global desktop OS market share,

underscoring its widespread adoption in the enterprise IT environment [22]. Similarly, Windows Server OS takes up a substantial market share of approximately 65% in 2022 [18]. On the other hand, the use of Windows OS in OT environments, which encompass industrial control systems (ICS) and critical infrastructure, is also significant. Many OT systems, such as Supervisory Control and Data Acquisition (SCADA) and Human-Machine Interface (HMI) applications, run on Windows-based platforms. The reference architecture developed by ATT&CK sets out Windows-based systems that are used as application server, engineering workstation, transient cyber asset (TCA), safety engineering workstation, etc. in OT environments [9]. Due to compatibility requirements with existing industrial hardware and software, several OT environments continue to operate on legacy Windows systems, such as Windows XP and Windows 7.

A recent incident with the CrowdStrike software update highlights a critical dependency on Windows-based systems. 8.5 million devices are affected, causing widespread disruptions in multiple industries worldwide [24]. Due to their widespread use and known vulnerabilities, Windows systems are often the primary targets of malware and ransomware attacks. In addition, Windows systems in OT environments may not always be configured with stringent security measures, potentially leading to unauthorized access and system compromises [15]. Reliance on legacy Windows systems in OT settings poses significant security challenges and requires proactive measures to mitigate associated risks.

The rise in ransomware attacks over the past few years further underscores the urgent need for advanced defensive mechanisms. With 93% of the ransomware targets Windows-based systems, enterprises require robust countermeasures to detect, analyze, and mitigate threats effectively [21]. Typically, honeypot systems are deployed by enterprise security teams, IT departments, government agencies, threat intelligence companies, cloud service providers, security researchers, and academics as an effective cyber security strategy for this purpose. They provide valuable threat intelligence, improve defensive capabilities, and serve as an early warning system against cyber threats. A variety of honeypot systems have been proposed and developed. Notable implementations include Cowrie [3], Glutton [4], OpenCanary [7], Conpot [2], T-Pot [10], AIDE [5], ICSNet [19], SIPHON [12] etc. However, despite previous work such as [16, 23], there has been very little state-of-the-art

*This work was done while the author was at iTrust, Singapore University of Technology and Design.

[†]Both authors contributed equally.

research, HopLab in [20] for instance, on Windows-based honeypot systems, which has created a significant research gap. On the other hand, conventional security solutions focus primarily on detection and response. However, a proactive approach involving deception could significantly enhance an organization’s security posture [8]. The need for additional security layers that do not depend solely on conventional endpoint protection solutions has never been more apparent [13].

To address these concerns, this paper introduces **HoneyWin** a high-interaction Windows honeypot that mimics an enterprise IT environment. The **HoneyWin** incorporates Windows 11 endpoints and an enterprise-grade firewall provisioning with network traffic capture, host-based event logging, deceptive tokens, endpoint security (EDR) and real-time intrusion detection alerts capabilities. Based on our review of existing state-of-the-art works, **HoneyWin** is a first-of-its-kind high-interaction Windows honeypot that has been designed, implemented, validated and deployed live in the wild.

HoneyWin offers the following contributions:

- **A scalable high-interaction Windows honeypot design.** The proposed **HoneyWin** system consists of three Windows 11 endpoints and an enterprise-grade firewall. However, **HoneyWin** was designed for scalability in the first place to allow us to conveniently expand it in the future. In particular, **HoneyWin** incorporates a lightweight container-based approach to establish public-facing honeypot devices. Both incoming and outgoing network traffic is captured separately with dedicated systems, not within the honeypot devices. Host event logs are not stored within the honeypot devices; instead, they are shipped directly to avoid detection and tamper proofing.
- **Holistic detection capability.** The **HoneyWin** captures both network traffic and host event logs. The honeypot devices are equipped with a commercial state-of-the-art endpoint security solution. By this means, **HoneyWin** accommodates the correlation between network traffic and host event logs and facilitates a holistic detection capability.
- **Deceptive tokens implementation and deployment.** To mislead and deceive adversaries who gain access to the honeypot devices, the **HoneyWin** implementation incorporates two types of deceptive tokens: (1) Spoofed Windows commands – commonly used Windows discovery commands that have been reconfigured to deceive attackers into the realm of navigating a legitimate enterprise network environment and (2) realistic bait files to further enhance the deception.

Organization: The remainder of this paper is organized as follows. Section 2 introduces **HoneyWin** beginning with a threat model and various design considerations, followed by the implementation of **HoneyWin** in an enterprise environment. Section 3 discusses penetration testing and malware detection testing to validate the implementation of **HoneyWin**. We provide experimental results from live deployment of the **HoneyWin** system in the wild for 34 days, including an in-depth analysis on the attacks received. Section 5 discusses the insights from the design and implementation of the **HoneyWin** system and the experimental results. Finally,

we provide related work in Section 6 and conclude the paper in Section 7.

2 HONEYWIN: HIGH INTERACTION WINDOWS HONEYPOT IN AN ENTERPRISE ENVIRONMENT

This section describes the proposed **HoneyWin** system. Firstly, we present the threat model and considerations that have been taken into account when designing the **HoneyWin**. Subsequently, we provide details of the implementation of **HoneyWin**.

2.1 Threat Model

We assume that the attacker has access to the honeypot by scanning the Internet or using a search engine such as Shodan¹. Once the honeypot is selected as the target, the attacker initiates a reconnaissance using port scan tools such as nmap² and identifies the exposed services (e.g., RDP port 3389). The attacker then probes each service to obtain more information and attempts to authenticate with a username and password. In this case, the attacker may brute-force or use other possible means to have the correct credentials. Upon gaining access, the attacker advances with the next phase of the cyber kill chain. The honeypot is designed to log the attacker’s interactions via network traffic and host event logs. In addition, the honeypot is implemented with a real-time alerting mechanism to report critical events, such as successful logins.

2.2 Design

This section discusses various considerations taken into account when designing a Windows-based honeypot in an enterprise environment.

High-Interaction Honeypot with Real Systems: We anticipate that the proposed honeypot setup includes real systems instead of low-interaction implementation that emulated certain services (e.g., Cowrie SSH/Telnet honeypot [3]). Having real systems as high-interaction honeypots maximizes the attack surface and allows full access to the underlying system. The setup features Windows 11 endpoints that are accessible directly on the Internet. Moreover, the setup incorporates an enterprise-grade gateway/firewall with Windows endpoints connected on the local area network (LAN), while the wide area network (WAN) side of the gateway is exposed on the Internet mimicking a typical enterprise environment.

Exposing Honeypots on the Internet: It is imperative that Windows endpoints and the enterprise gateway require public IP addresses to make them accessible on the Internet. Windows endpoints and the gateway could be hosted with Infrastructure as a Service (IaaS) or Virtual Private Server (VPS) which belong to cloud service providers (e.g. Amazon Web Services). Using Virtual Private Network (VPN) service provides public IP addresses from servers located in more diverse geolocations while the devices could remain within the perimeter of our setup.

Network Traffic Collection and Host-based Event Logging: To detect and analyze intrusions and threats received by devices in our honeypot, it is essential to capture all incoming network traffic. In addition, we must capture the outgoing network traffic initiated

¹<https://www.shodan.io/>

²<https://nmap.org/>

from the devices, as these may be attempts by malicious actors to establish network connections to command & control (C2) servers, scanning and targeting other vulnerable devices elsewhere. Modern operating systems (e.g., Windows, Linux) and security appliances (e.g., Cisco, Fortinet) feature logging of events related to the system, security, and applications. Host logs provide fine-grained visibility into the activities that occur in each device, allowing better detection, investigation, and response that complements network traffic analysis.

Incorporation of Deceptive Tokens: Deceptive tokens are crucial to making the honeypot a realistic high-interaction environment providing a high level of interactivity with the system while preventing attackers from actually affecting the system and network. The deceptive tokens in the form of windows built-in executables, would provide benefits such as being easily configurable, consistent, and scalable across multiple honeypot deployments. The spoofed windows commands reading from a common configuration file that is easily replicable allow for quick deployment across multiple honeypot devices. In addition, a modular design approach allows partial modifications without affecting the entire configuration and making updates seamless. To deceive the attackers, these commands would be reconfigured so that they deceive the attackers into believing that they are navigating a real corporate network environment. In the event that an attacker gains access via remote means, deceptive tokens ensure that the system presents misleading but realistic network and system information. By designing the tokens in the form of built-in system commands, these modified executables generate and return false, yet plausible system responses, reinforcing the illusion of a convincing corporate network. This approach enhances adversary engagement while allowing defenders to gather valuable intelligence on the behavior of the attacker.

Real-time Intrusion Detection and Alerts: Having Windows endpoints and the gateway accessible from anywhere implies a rich attack surface. Moreover, there is a significant security risk once the attacker compromises and takes control of the systems. To mitigate such risk and take appropriate actions, the setup should implement real-time intrusion detection (e.g., successful logins) and alerting mechanism.

Network Traffic and Log Backup: Through the reconnaissance, the attacker gains access to the Windows endpoints and gateway. Therefore, network traffic collection shall not take place within the systems in the first place. Similarly, the host logs shall not be kept within the systems; instead, they should be forwarded and stored in a security-hardened system. Even in such a scenario, it is still possible that either the collected network traffic or host logs are compromised. To minimize the risk of losing such invaluable data, the set-up shall incorporate a backup system to store the replica of captured data and logs. The backup system shall not have an Internet connection and has access to the systems storing the data and logs but not the other way around.

Cloning and Restoring Windows Systems: It is necessary to save the system image before they are deployed live on the Internet since it is likely that the systems get compromised over time and may become out of control. Once sufficient data are captured to attribute the attack tactics, techniques and procedures (TTPs), the

systems could be restored into their pre-deployment state. Depending on the attack vectors, certain adjustments to the system image may be necessary before redeploying live on the Internet.

Log Management and Analysis Platform: To efficiently manage and analyze large volumes of log data, it is essential to incorporate a log management and analysis platform such as ELK³ or Splunk⁴ into our setup. These platforms are capable of collecting, parsing and storing large amounts of data, as well as searching, filtering, analyzing, and visualizing the collected data.

Automation and Orchestration: It is expected that the honeypot is deployed live around the clock. Hence, various design considerations discussed above such as setting up of VPN connections, network traffic and host logs collection, intrusion detection and alerts, etc. shall be automated. In addition, monitoring of the status of the entire setup and orchestration of various components independently are required.

2.3 Implementation

Taking into account the design considerations discussed in Section 2.2, we envision and propose high-interaction Windows honeypot system in an enterprise environment, namely **HoneyWin**. Figure 1 shows the proposed implementation of **HoneyWin**. The setup consists of three Windows 11 endpoints (E1, E2 and E3) and an enterprise gateway (G1) as honeypots. The endpoint E3 and the gateway G1 are directly accessible on the Internet, while the endpoints E1 and E2 reside within the private network (i.e., the 'Enterprise Network' (EPN) in Figure 1). They are connected to the gateway and have Internet connection via Network Address Translation (NAT) scheme. The incoming connection therefore must bypass the gateway to access the E1 and E2 endpoints.

VPN Forwarding with Docker Container: To expose devices on the Internet, we adopted a lightweight approach to establish secure tunnels to VPN servers and acquire public IP addresses. In particular, a customized docker container has been developed with the functionalities to establish VPN tunnels to servers in specified geolocations, set up port forwarding to the devices on the 'Honeypot Network' and automatically capture the incoming network traffic received on exposed public IP addresses. Typically, each device (i.e., E3 and G1) has a dedicated docker container. However, it should be noted that our implementation allows multiple containers to forward the network traffic to the same device, thereby increasing the geographic presence given a limited number of real systems in case. The containers run inside the Workstation (U0). To automate the number of containers and their names, public IP addresses, open ports, and forwarding IP addresses to the devices on the 'Honeypot Network' (HPN), a set of shell scripts and a CSV file are used. To expose a specific device on the Internet, the user first connects the device to the HPN, then updates the CSV file with the necessary information, and uses the script to establish end-to-end connectivity between the VPN server and the device. With the aid of a shell script and also thanks to a docker container, we could bring all containers or individual ones online or offline

³<https://www.elastic.co/elastic-stack>

⁴<https://www.splunk.com/>

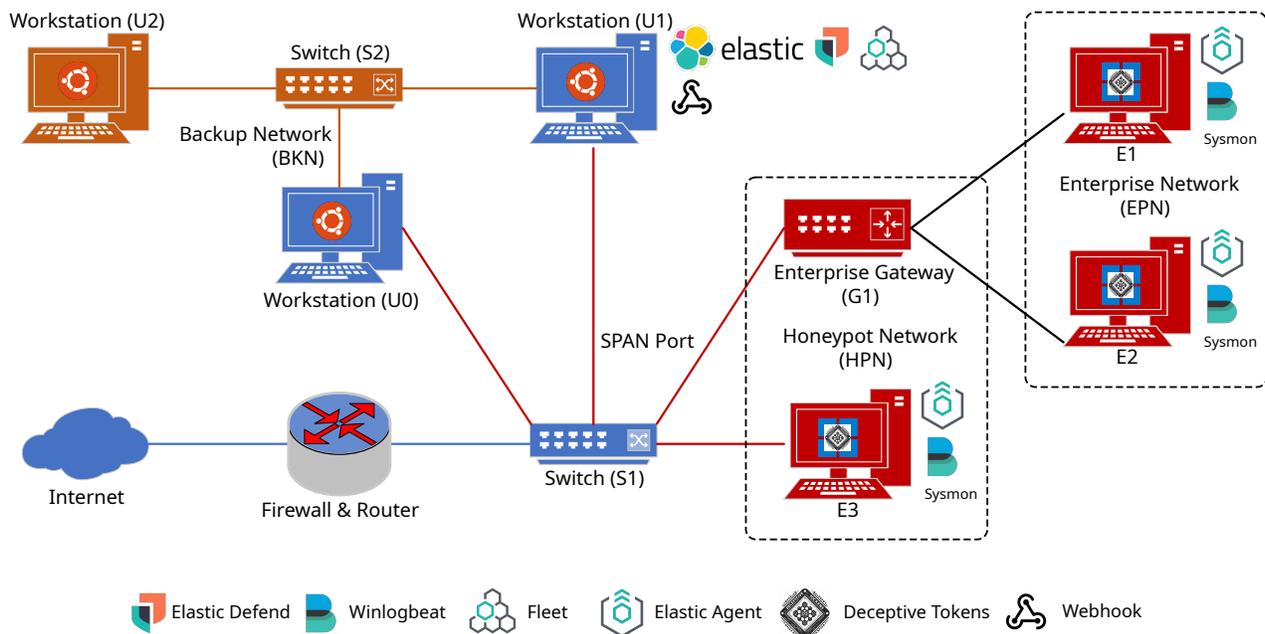


Figure 1: Windows-based Honeypots in Enterprise Environment

within seconds. In addition, the IP addresses of the exposed devices are mapped to a registered domain to mimic an enterprise environment.

Incoming and Outgoing Network Traffic Collection: As described in the previous section, incoming network traffic to public-facing devices is captured in each docker container. Once the attacker is within the HPN or the EPN, outbound connections may be initiated to contact C2 servers, scan and target other vulnerable devices elsewhere since all devices on both networks have the Internet connection. We use the SPAN port of the switch (S1) to capture the outgoing network traffic and store the PCAPs in the Workstation (U1).

Host Log Collection and Monitoring with ELK Stack: We also collect host logs from the honeypot devices. For all Windows endpoints, the System Monitor (Sysmon) is installed to monitor and log system activity to the Windows event log. Sysmon⁵ provides detailed information on process creations, network connections, and changes to file creation time. Sysmon runs as a protected process and does not allow for a wide range of user-mode interactions. However, Sysmon does not provide analysis of the events and does not attempt to hide itself from attackers. As such, we rename the Sysmon executable and the driver to hide it and add complexity to adversaries attempting to identify security tools in use as a first line of defense. For the enterprise gateway, we enable Syslog logging via the administrator dashboard.

Instead of storing the host logs within the honeypot devices, we have setup an ELK stack in Workstation (U1) for log management and analysis. In this case, an Elastic Agent (EA) is installed in each Windows-based system. For the gateway (G1), an intermediate system is setup and the Elastic Agent is installed to collect and ship

the Syslog to U1. The Elastic Agent⁶ provides a unified way to add monitoring for logs, metrics, and other types of data to a host device. To prevent the EA from being uninstalled without authorization, agent tamper protection is also enabled in the EA policy. We have also set up Fleet⁷ in the ELK stack and installed the Fleet Server in U1 to manage EAs and their policies. Elastic provides integrations, which are prepackaged assets allowing users to collect, store, and visualize data from various sources with ELK. Elastic Defend⁸ is one of the available integrations in ELK to detect, prevent, and respond to cyber threats. We have installed Elastic Defend, Prebuilt Security Detection Rules, Windows, and enterprise gateway logs integrations with our setup. Prebuilt Security Detection Rules integration stores the security rules to detect malware in the endpoints and generate alerts. Windows integration is to receive the Sysmon logs forwarded by the endpoint EA whereas enterprise gateway logs integration is to receive the Syslog from the intermediate system of the gateway, respectively.

Successful Login Alerts: Based on our threat model in Section 2.1, the attacker attempts to log in to our honeypot devices through the exposed services. In our setup, we have opened RDP (i.e. Port 3389) and SSH (i.e. Port 22) for Windows endpoints, whereas HTTPS (i.e. Port 443) is open for access to the administrator dashboard of the enterprise gateway. The attacker who is able to successfully log in with correct credentials signifies a breach to our devices. Hence, it is critical to provide real-time notifications of successful login to any of our devices. To realize this capability, we have installed Winlogbeat in our Windows endpoints and configured it to capture successful RDP and SSH login events, and forward the

⁵<https://learn.microsoft.com/en-us/sysinternals/downloads/sysmon>

⁶<https://www.elastic.co/elastic-agent>

⁷<https://www.elastic.co/guide/en/fleet/current/fleet-overview.html>

⁸<https://www.elastic.co/guide/en/integrations/current/endpoint.html>

Table 1: Deceptive Tokens, Installation Paths and Hostnames

Deceptive Token	Path	Hostname
wmiutils.mof ipconfig.exe net.exe netstat.exe arp.exe systeminfo.exe	C:\Windows\System32\wbem	E1, E2, E3
Current-5G-Network-Infrastructure-PPP-in-H2020_Final_November_2023.pdf 5G_vendor_technology_list.xlsx TDoc_List_meeting_SA#91-e.xlsx Security-in-5G.pdf	C:\Users\<E3>\Desktop	E3
db_info.conf iam_access.conf	C:\Users\<E3>\Documents	
aws_creds-list_14082024.xlsx intelsat-mobile-live-5G-infra.pdf	C:\Users\<E2>\Documents	
ESEC Structure.png huawei-5g-cpe-pro-custom.pdf IC-Project-Team-Meeting-Minutes-11856.xlsx	C:\Users\<E2>\Desktop	E2

event logs to the Logstash component installed in U1. We then use a webhook integration to send out emails to notify the successful login events.

In addition, we have also configured the Winlogbeat to capture failed logins, which indicate brute-force attempts. The events are forwarded and stored in U1. For the case of enterprise gateway, we have created an automation with ‘Admin login successful’ event as a trigger. Upon activation, the webhook sends an email alert.

User Accounts and Login Credentials: For each Windows endpoint, we have one user account with administrative privileges and one or more standard users. The password for the administrative user is a randomly generated one with 16 characters consisting of uppercase, lowercase, digits, and brackets. This has been the design decision to prevent the adversary from taking complete control of the endpoint. However, the passwords for standard users are weak passwords with just eight characters consisting of uppercase, lowercase, and digits. For the gateway (G1), there is only one administrative user and the password is set as a random generated one with 16 characters consisting of uppercase, lowercase, digits, and brackets.

Firewall and Enterprise Gateway: In Figure 1, we have set up a ‘Firewall & Router’ to manage VLAN trunks, DHCP reservations, DNS settings, firewall rules, etc. For example, the outgoing connections from the HPN are blocked by default. Selected devices are given Internet access through firewall rules. Similarly, the enterprise gateway/firewall (G1) is also configured to have fine-grained control and access management of devices in the EPN.

Backup Server and Clonezilla: We have setup Workstation (U2) as a backup server to store the incoming network traffic captured in U0, the outgoing network traffic captured via the SPAN port with U1, successful logins and failed login attempts, and Elasticsearch indices of Sysmon events. The U2 connects to U0 and U1 on the ‘Backup Network’ (BKN) and does not have the Internet

connection. Also, only U2 has access to U0 and U1 but not vice versa. We use Clonezilla to keep the system image of Windows endpoints so as to restore into the pre-deployment state if the system gets compromised and becomes out-of-control. In addition, we save the configurations of ‘Firewall & Router’ and enterprise gateway (G1).

Scalable Design: While the current **HoneyWin** setup consists of a limited number of honeypot devices: an enterprise gateway and three Windows endpoints, we have designed it for scalability which enables us to expand it conveniently in the future. Having a light-weight container-based approach to establish secure VPN tunnels, we could extend the setup with additional public-facing endpoints. Moreover, we could place and interchange the honeypot devices at different geolocations during specific time periods. Capturing of incoming and outgoing network traffic inside each container and U1 respectively, real-time alerting of successful logins to the honeypot devices, backing up of PCAPs and Elasticsearch indices are fully automated. The state-of-the-art ELK stack with Elastic Defend, Prebuilt Security Detection Rules, Windows and Firewall Logs integrations, as well as Fleet and Elastic Agent allows us to secure the endpoints from tampering, detect, and analyze TTPs.

Deceptive Tokens: We have also incorporated deceptive tokens into the Windows endpoints. Table 1 shows the deceptive tokens allocation and installation in three Windows endpoints. The objective is to mislead adversaries who gain access to the system by manipulating the output of commonly used Windows discovery commands, such as systeminfo, ipconfig, netstat, net.

‘systeminfo’ is commonly used for host enumeration by attackers as it provides detailed system information, useful for attackers to plan their next course of action, as key details include windows version, build number and installed patches along with domain membership and hostname as well as other useful data. ‘ipconfig’ is used for network reconnaissance as it provides critical networking information that could be used for lateral movement.

It could also be used to list Domain Name System (DNS) and Dynamic Host Configuration Protocol (DHCP) servers that could be further leveraged by attackers to perform other network based attacks such as DNS poisoning. ‘netstat’ is used for network reconnaissance as it lists active remote connections in the host system and provides a list of connected hosts that could be pivoted to by the attacker. ‘net’ allows adversaries to list user accounts on the local system as well as domain users and their group members, giving crucial data for privilege escalation on systems or allowing attackers to install a backdoor by using the binary to add or remove users from the system. Additionally, it allows the enumeration of network shares for data exfiltration or lateral movement, making the executable highly valued for living off the land movement. The deceptive tokens listed in Table 1 were specifically selected and created to replace built-in windows executables that have been identified to be key binaries exploited by adversaries/nation state actors during their operations where part of their unique trade craft involved.

After the design and creation of the deceptive tokens, implementation in the **HoneyWin** system requires careful execution to prevent attackers from detecting legitimate binaries. The deceptive tokens and bait are distributed in various locations on the different **HoneyWin** systems, as provided in Table 1. ‘ipconfig’, ‘net’, ‘netstat’, ‘arp’ and ‘systeminfo’ were placed on all three hosts on ‘C:\Windows\System32\wbem’ together with ‘wmiutil.mof’ which contains the configuration data used by the different binaries. The ‘wmiutils.mof’ file is placed in the Web-Based Enterprise Management folder to hide among other Managed Object Format files that are used by the Windows Management Instrumentation executable. The altered binaries were placed there to enhance the legitimacy of existing within the ‘System32’ folder and are added to the user path from there to take precedence over the original binaries, which are left in their original locations for future use.

To further enhance the deception, the team has deployed realistic bait files along with these deceptive executables. These files include configuration files, project meeting minutes, network infrastructure reports, and AWS Identity and Access Management (IAM) credentials that appear to provide access to cloud services. By planting such high-value artifacts, the system attracts attackers to interact with the deceptive environment, thereby increasing the likelihood of engagement and intelligence collection. It should be noted in Table 1 that <E2> and <E3> are placeholders and represent standard users for the endpoints. As mentioned earlier, each endpoint may have more than one standard user.

3 VALIDATION OF HONEYWIN

The primary objective of evaluating the **HoneyWin** implementation is to confirm its effectiveness in mimicking a realistic enterprise environment and capturing real-world attack behaviors. In particular, the validation process has been designed to measure how effectively the honeypot detects and responds to network reconnaissance, unauthorized access attempts, and privilege escalation, while also determining its overall detection visibility, response time, and ability to capture all intended data.

By simulating common adversarial tactics, ranging from port scans and lateral movements to sophisticated privilege escalation

methods, the validation effort provided detailed insight into monitoring granularity, generating alerts, and the ability to withstand more advanced attack strategies targeting **HoneyWin**. This rigorous testing not only validates the design assumptions of **HoneyWin** but also highlights areas for further refinement, ensuring that it provides high fidelity in both deception and detection capabilities.

3.1 Penetration Testing

Penetration testing formed the cornerstone of the validation process for **HoneyWin**, ensuring that its network monitoring and logging systems are functional and collect the desired data. Before incorporating the deceptive tokens, the environment is rigorously tested to confirm proper implementation.

Network Reconnaissance: A series of ‘nmap’ scans are simulated adversarial port-scanning behavior. These scans revealed limited exposure and strong initial defenses:

- **Top 1000 and all available ports:** Port 514 was consistently displayed as filtered, demonstrating restrictive policies typical of secure environments.
- **Targeted scans on ports 22 (SSH) and 3389 (RDP):** Both returned “host down” statuses, underscoring effective blocking of unnecessary exposure.

Host-Level Testing: Tools and commands such as ‘netstat’, ‘ipconfig’, ‘route print’, and ‘arp’ were executed to evaluate network configurations and logging capabilities of **HoneyWin**. The results demonstrated that logging configurations effectively track host-specific actions.

Windows Privilege Escalation Awesome Scripts (WinPEAS)⁹ is a post-exploitation tool developed to help security professionals, penetration testers and ethical hackers identify opportunities for escalation of privileges in Windows systems. Our test results show that the default Windows Defender in Windows 11 Professional is able to thwart WinPEAS. Similarly, commands seeking to access sensitive directories and files, such as the Security Accounts Manager (SAM) file and tasklist, were blocked, proving that endpoint defenses provided robust real-time security.

Privilege Escalation Attempts: Mimikatz¹⁰, a powerful post-exploitation tool that allows users to view and save authentication credentials such as Kerberos tickets, is used in privilege escalation tests for credential dumping. These attempts are effectively mitigated by Microsoft Credential Guard, which prevented access to critical credential storage. The usage also ensures that there is visibility of host-based actions and executables allowing us to monitor even down to what dynamic link libraries (DLLs) are called.

Access to ELK Stack: During the testing process, we identified misconfigurations during the scans where the ELK stack opens ephemeral ports and allows unrestricted access. An initial access to a workstation in the EPN led to the discovery of unintended access to an internally hosted ELK stack in U1, which is caused by the use of an earlier ELK version (e.g. Elasticsearch 7.17) for which security is not enabled automatically during the installation. In particular, the ELK stack does not require authentication, allowing unrestricted access to the Elasticsearch database and the Kibana dashboard. The access effectively provides insight into the defensive monitoring

⁹<https://github.com/peass-ng/PEASS-ng/tree/master/winPEAS>

¹⁰<https://github.com/ParrotSec/mimikatz>

capabilities. The ELK stack is reinstalled fully with security features and confirms that access required authentication.

3.2 Testing Detection and Alerts with C2 Malware

As part of the validation process, we have also tested the **HoneyWin** with 2 sets of malware on the ELK stack with the Elastic Defend setup. One malware is a custom designed malware that provides a reverse shell via the 'netcat' connection. The malware provides a reverse shell connection, which allows an adversary to remotely control the endpoint through the command line interface (CLI). We created this malware as a custom shellcode which connects to the adversary IP address and port while providing a CLI to interact with the endpoint upon execution.

Another malware was a modified Havoc malware. Havoc¹¹ is a post-exploitation and C2 framework designed for red-teaming and adversary simulation. We created a Havoc agent with customized settings for stealth and evasive properties. In this case, the agent runs on the endpoint and establishes a connection to the adversary's Havoc listener, enabling the adversary to remotely control the endpoint by sending commands. Both malware were designed to be able to bypass Windows Defender, but intended to be detected by Elastic Defend. Elastic Defend is able to detect both malware through their prebuilt Malware Detection Alerts rules.

4 EXPERIMENTAL RESULTS

We have deployed the **HoneyWin** system live in the wild for 34 days from 22 August 2024 to 25 September 2024. This section provides our analysis on incoming and outgoing network traffic, failed login attempts, successful breaches via the exposed services, and the attacks initiated by the adversaries.

4.1 Network Traffic Analysis

During this operating period, **HoneyWin** received ~5.79 million unsolicited connections, of which the gateway received 2.22 million connections and the Windows endpoint (E3) received 3.57 million connections. In general, the Windows endpoint (E3) received 60% more connections than the gateway (G1). Figure 2 shows the distribution of the connections received by G1 and E3 over a 12-hour interval for 34 days. G1 received more connections than E3 for 8 intervals only.

Figures 2b and 2c show geolocations based on IP address lookup of incoming connections received by G1 and E3, respectively. For the gateway (G1), 38% and 13% of the connections are initiated from IP addresses belonging to United States and Philippines. For the Windows endpoint (E3), 43% and 19% of the connections are from IP addresses belonging to Russia and the United States. It should be noted that the identification of geolocations from IP address lookup is only indicative since the attackers may have spoofed their true IP addresses.

Figures 3a and 3b show connections to open port of the gateway (G1) and Windows endpoint (E3) over 12-hour intervals for 34 days. The vertical axis in both figures shows the number of sessions in logarithmic scale. The gateway (G1) receives more than 2,123

HTTPS sessions per 12 hours on average. On the other hand, the Windows endpoint (E3) receives more than 31,972 RDP and SSH sessions combined every 12 hours on average. This indicates that the Windows endpoint (E3) is more than 15 times active compared to the gateway (G1). Furthermore, the Windows endpoint (E3) receives 59,591 RDP and 4,263 SSH sessions per 12 hours on average. The attacker actively exploits RDP nearly 14 times more than SSH.

4.2 Failed Login Attempts

As described in Section 2.3, **HoneyWin** is capable of capturing failed login attempts for the Windows endpoint (E3). Figure 4 shows a boxplot of failed login attempts for E3. The vertical axis shows the number of failed login attempts per hour for each day on the horizontal axis. The data range varies over the **HoneyWin** operating period. We observe significant outliers in certain days (e.g., Sep 01, 02 and 05). Also, significantly higher first-quartile (Q1) and median (Q2) values (e.g., Sep 10 and 11). This strongly indicates active brute-forcing attempts by the attackers during certain hours of the day or throughout certain days.

4.3 Successful Logins

During the operating period, the **HoneyWin** detected 5 successful logins via RDP and 354 successful logins via SSH. We have analyzed all successful log-ins and developed an attack attribution method which correlates the incoming network traffic, host logs, and outgoing traffic holistically. For 4 out of 5 successful logins via RDP, we find that these logins are short-lived and do not see any activities performed by the attackers after the successful logins. Section 4.4 provides our analysis on the remaining successful login session during which the attacker interacted with our deceptive tokens. As described in Section 4.1, although the attacker actively exploits RDP nearly 14 times more than SSH, our analysis shows that a stealthy attack was launched through successful SSH logins. We provide detailed analysis of this attack in Section 4.5.

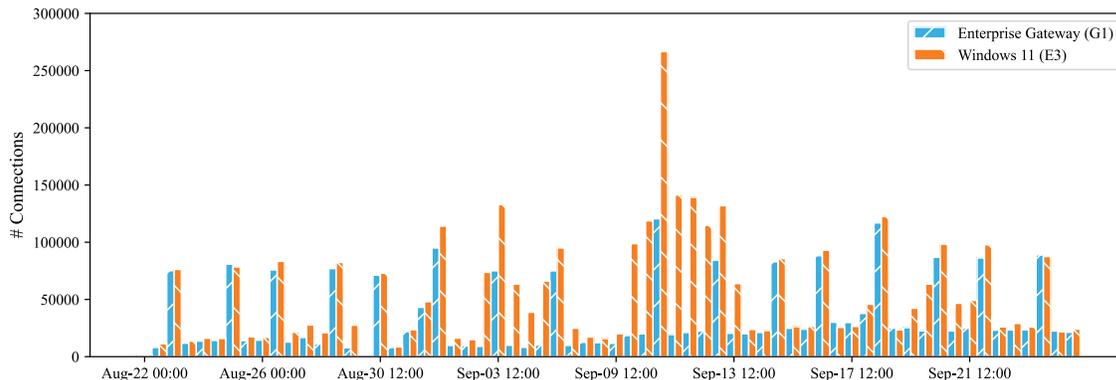
However, no successful login alerts are generated for the gateway (G1) and two Windows endpoints (E1 and E2). We conjecture that strong password setting of the gateway (G1) prevented the breach, although there are sustained brute-force attempts. For the case of E1 and E2, the endpoints are on the EPN with NAT connections to the G1. This imposes an additional layer of difficulty that requires one to bypass the gateway G1 to reach both endpoints.

4.4 Analysis on Deceptive Token Interaction

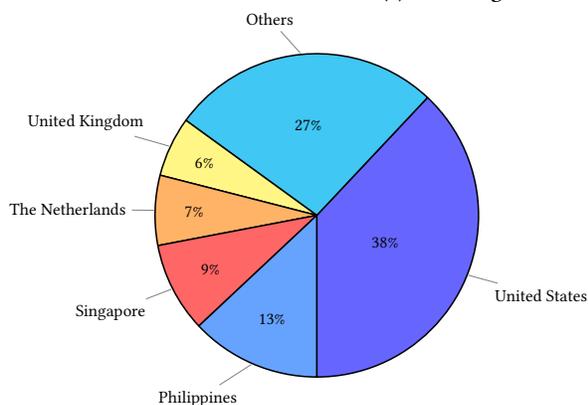
Analysis of system logs and network traffic reveals that a successful RDP session was established using one of the standard user accounts of E3. Correlating process creation logs with network traffic, we are able to trace the attacker's activities and interactions with our deceptive tokens.

Based on the timestamp of the successful RDP login alert notification, we identify the RDP login event and correlate it with the process creation timestamps. Typically, when a user logs in via RDP, the system spawns key processes such as 'smss.exe', 'winlogon.exe', and 'userinit.exe'. The sequential execution of these processes confirmed a successful login as illustrated in Figure 5.

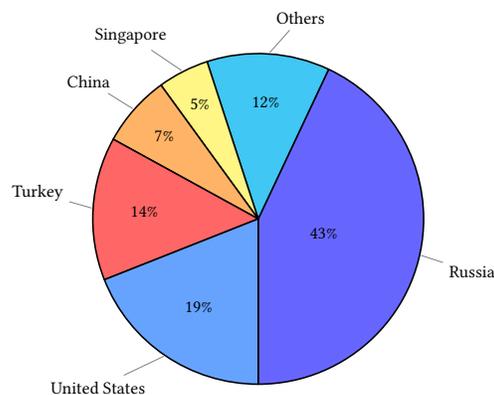
¹¹<https://github.com/HavocFramework/Havoc>



(a) Incoming Connections Received by G1 and E3

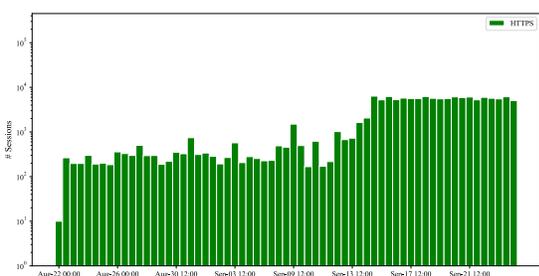


(b) Geolocations of Incoming Connections Received by G1

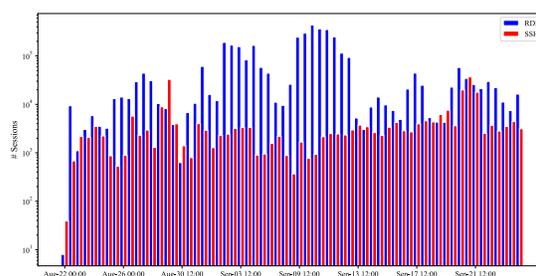


(c) Geolocations of Incoming Connections Received by E3

Figure 2: Analysis of Incoming Connections Received by G1 and E3



(a) HTTPS Sessions of Enterprise Gateway (G1)



(b) RDP & SSH Sessions of Windows 11 (E3)

Figure 3: Connections to Open Ports of G1 and E3

Further analysis of the process creation sequence reveals that 'powershell.exe' was launched by 'explorer.exe', indicating post-login activity. Notably, there is a two-minute interval between the login and the execution of powershell.exe. This suggests that the reconnaissance is performed manually by a user rather than

through an automated script, reinforcing the likelihood of an active attacker on the keyboard.

A key observation is that 'powershell.exe' subsequently executed our deceptive token, 'systeminfo.exe'. This confirms that

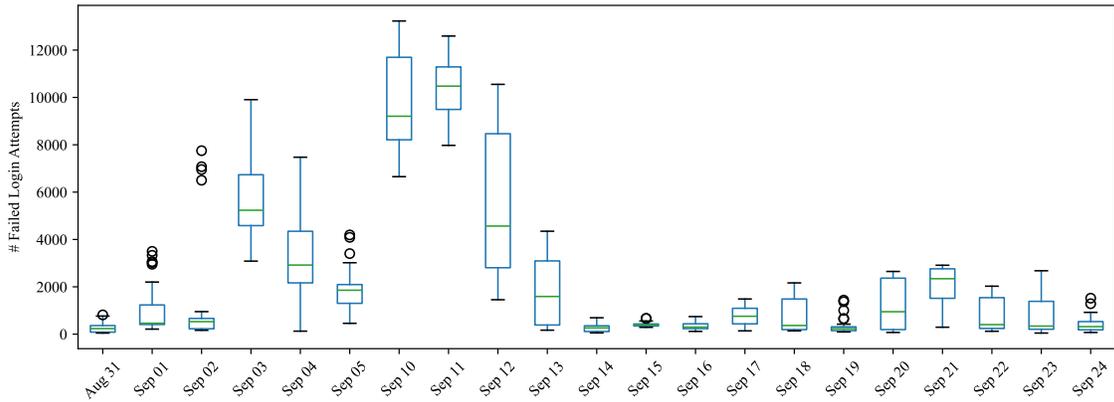


Figure 4: Boxplot of Failed Login Attempts

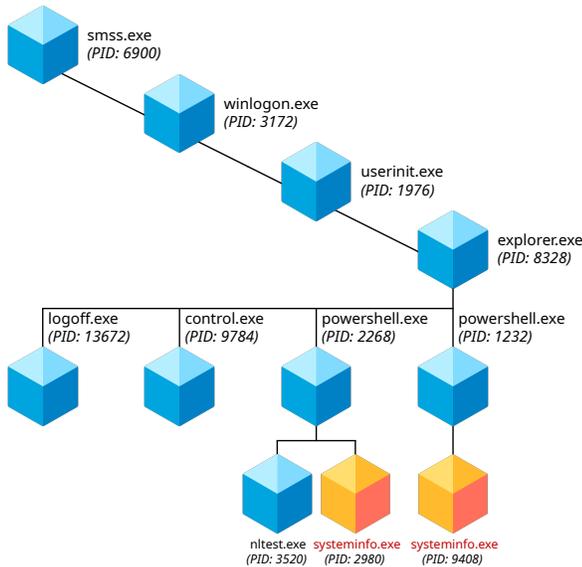


Figure 5: Analysis of a Successful RDP Login & Deceptive Token Interaction by the Adversary

the attacker has been engaging with the deceptive environment, believing it to be a legitimate endpoint.

4.5 Stealthy SMTP Brute-Force Bot Attack via Successful SSH Logins

We received a successful SSH login alert at 08:40 PM on 18 September 2024 (GMT), followed by two hourly successful logins until 04:50 PM, 23 September 2024. Our initial analysis on these successful logins does find any sign of attack. Subsequently, the successful logins become every two minutes. Although we did not find significant host activities up to this time, the SPAN port is capturing 100 MB of network traffic every 7 minutes. Further investigation shows reverse SSH tunneling after the initial successful SSH login

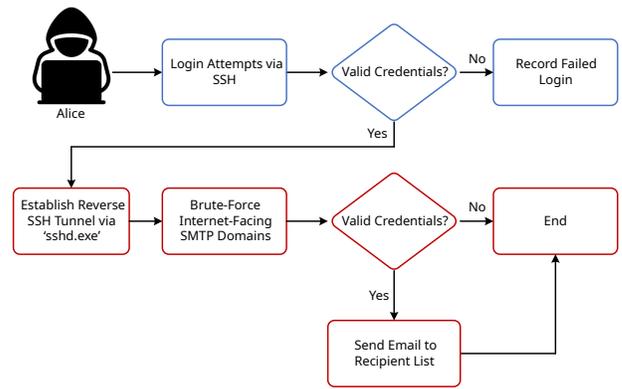
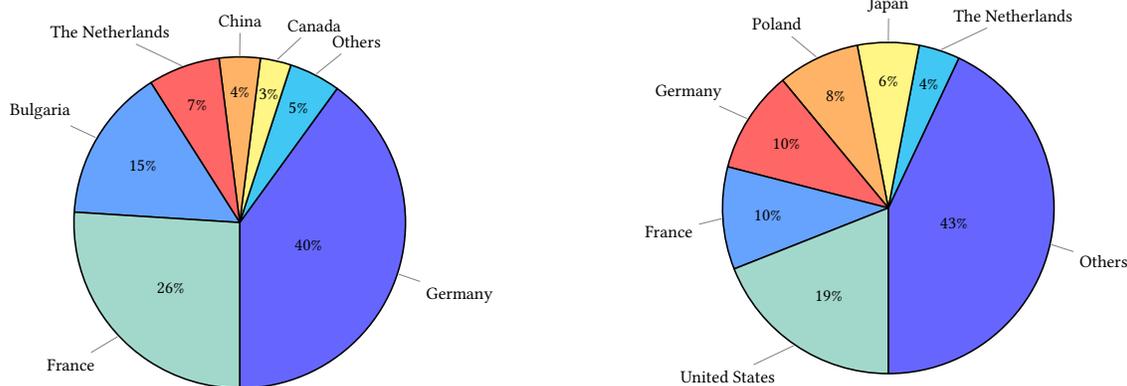


Figure 6: Analysis of Successful SSH Logins & SMTP Brute-Force Bot Attack

spawning child SSH processes. In this case, SMTP outgoing traffic could be observed from child SSH processes. However, the ELK is not able to log the SMTP payload due to reverse SSH tunneling crippling host-based detection and insights. Since all outgoing traffic originating from the **HoneyWin** is captured via the SPAN port, we are able to identify the SMTP payloads from the network traffic. The adversary was performing an attack on global SMTP domains via successful SSH logins recruiting the endpoint (E3) into a botnet. Figure 6 shows our analysis of successful SSH Logins followed by SMTP brute-force bot attack.

It is evident that host-based logging alone does not suffice and shed light on the advantage of **HoneyWin** as it captures incoming network traffic, host logs and outgoing network traffic, all together thereby facilitating a holistic detection capability. We have developed an attack attribution algorithm to establish an end-to-end correlation of host logs and network traffic (see Algorithm 1).

We are able to trace the attacker's activities and determine that SMTP brute force bot attacks have been carried out on Internet-facing SMTP servers. Figure 7a shows the geolocations of 354 successful SSH logins based on IP address lookup of incoming network



(a) Geolocations of Incoming Connections for Successful SSH Logins (b) Geolocations of Outgoing Connections for Successful SSH Logins

Figure 7: Analysis of Incoming and Outgoing Connections for Successful SSH Logins

Table 2: Top 15 SMTP Usernames

Username	# Occurrence
info	2529
contact	562
no-replay	483
admin	465
office	458
inf	389
marketing	326
mail	324
support	315
webmaster	314
hello	248
noreply	246
root	223
test	215
postmaster	214

Table 3: Top 15 SMTP Passwords

Password	# Occurrence
co23	587
in23	398
ma23	380
123456	241
q1w2e3r4	194
ne23	184
123654	180
test123	177
11223344	175
P@ssw0rd	175
1234	174
Abcd1234	172
1q2w3e4r	170
11111111	166
sa23	164

Table 4: Top 15 Username & Password Pairs

Username	Password	# Occurrence
info	in23	128
info	inf23	39
marketing	marketing@1	31
hello	Hello@15	28
no-replay	noreplay@1	27
test	test@1	27
webmaster	Webmaster@15	26
contact	Contact@1	25
contact	Contact@15	24
root	Root@1	24
contact	contact@15	24
contact	Contact2021	23
office	office@1	23
root	root@1	22
root	root@15	22

traffic connections, noting that attackers could have spoofed their true IP addresses. Similarly, Figure 7b shows geolocations of outgoing network traffic initiated by the attacker. In this case, geolocations may not have spoofed since these are the attacker’s targets. In fact, the outgoing traffic indicates geolocations of 180 countries while the United States taking up 19% followed by France and Germany with 10% each, Poland (8%), Japan (6%) and The Netherlands (4) where the remaining 43% belongs to 174 countries. In terms of outgoing port distribution, port 25 (Standard) takes up 74%, 587 (Default) and 465 (TLS) are 22% and 4% respectively.

Our analysis shows that the attacker used 151,179 credentials. Table 2 provides the top 15 SMTP usernames, while Table 3 and 4 list the top 15 passwords and username-password pairs used by the attacker. The brute-force attempts harvested 1250 successful SMTP credentials. The attacker uses the email addresses in Table 5 to record the successful credentials upon acceptance by the SMTP

server. The successful SMTP credentials do not appear to be randomly generated, rather they are specific to certain users. These credentials likely originated from a data breach, and we choose not to reveal them in this work.

5 DISCUSSION

This section discusses our insights from the design and implementation of the **HoneyWin** system and the experimental results. Firstly, three endpoints (i.e., E1, E2 and E3) deployed with the **HoneyWin** system are not virtual machines. They are rather real devices (e.g., Mini PC) running a fresh out-of-the-box (OOB) Windows 11 Professional with latest updates and automatic updates enabled. The deceptive tokens and bait files are installed on top of the OOB. As provided in Table 1, the deceptive tokens are installed system-wide for all users. However, bait files are placed for certain standard user accounts only recalling that <E2> and <E3> are placeholders for standard users. There are no bait files for the standard user account

Algorithm 1 Attack-Attribution-SSH

Input: Successful SSH login alerts **A**, Incoming traffic PCAPs **N**, SPAN port PCAPs **S**, Elasticsearch database **E**

Output: Incoming IPs **I**, Outgoing IPs & Ports **O**, Payload: **P**

- 1: Get timestamps **T** from alerts **A**
- 2: Construct Elasticsearch query **Q** for **T**
- 3: Construct TShark query **F** for **T**
- 4: $I = \text{GETINCOMINGATTACKERIP}(N, F)$
- 5: $O = \text{GETOUTGOINGDATAFROMELASTIC}(E, Q)$
- 6: Construct TShark query **G** to extract SMTP payload from **O**
- 7: $P = \text{GETOUTGOINGSMTPPAYLOAD}(S, G)$
- 8: **return I, O and P**

9: **procedure** `GETINCOMINGATTACKERIP(N, F)`

- 10: Get incoming attacker IPs **I** from PCAP **N** with query **F**
- 11: **return I**
- 12: **end procedure**

13: **procedure** `GETOUTGOINGDATAFROMELASTIC(E, Q)`

- 14: Initialize Elasticsearch connection
- 15: Get outgoing IPs & ports **O** from Elasticsearch **E** with **Q**
- 16: **return O**
- 17: **end procedure**

18: **procedure** `GETOUTGOINGSMTPPAYLOAD(S, G)`

- 19: Get outgoing SMTP data (username, password, mail to, mail from, etc) **P** from **S** with query **G**
- 20: **return P**
- 21: **end procedure**

Table 5: Attacker Email Addresses

Email Address	# Occurrence
toron@imobust.com	172
no-reply-1@mx-test-serv.org	155
no-reply10@mx-test-serv.org	147
c2@mail-master.org	136
bt@mail-master.org	126
c4@mail-master.org	101
c3@mail-master.org	98
c1@mail-master.org	86
check@bewareofdogs.xyz	85
no-reply-2@mx-test-serv.org	80
mail2@glob22glo1.su	41
no-reply12@mx-test-serv.org	26

breached for the E3 endpoint during the successful RDP session where the adversary interacted with the deceptive token. This explains the fact that the adversary is not able to interact with the bait files in this case and highlights the need for a scheme to distribute the bait files and the deceptive tokens automatically across all the endpoints.

Successful SSH log-ins for the SMTP brute-force bot attack occurred every two hours initially, indicating that it was likely a

scripted attack. The use of reverse SSH tunnels reveals the adversary’s intention to avoid host-based detection and maintain stealth as long as possible. Successful SSH logins may have been classified as false positives, and the attack could have been missed in the absence of outgoing network traffic capture.

While this work focuses on the HoneyWin deployment in an enterprise environment, we could extend **HoneyWin** to an OT environment with HMI and SCADA endpoints, potentially making it more enticing for the adversary.

6 RELATED WORK

Honeypots are security resources whose value lies in their ability to be probed, attacked and compromised. In general, there are two types of honeypots: first, low-interaction honeypots, which present simulated or emulated services/environments to attackers, and second, high-interaction honeypots, which show real systems to attackers. The characteristics of honeypots are that they are deceptive, discoverable, interactive, and monitored.

Several prior efforts have leveraged honeypot systems as a proactive mechanism for network defense, recognizing their potential to gather in-depth threat intelligence and lure attackers away from real assets. In [14], the authors highlighted the value of honeypots to test adversaries’ behaviors in controlled yet realistic environments, while other works have focused on enhancing honeypot scalability and adaptability in complex networks. As Grimes described, every corporate entity should be running honeypots if they are interested in the earliest warning possible of a successful hack or malware infiltration [11].

Various honeypot systems such as honeypots for database systems, web applications, services, ICS/SCADA, etc., have been proposed and developed by industry and academia [1]. Notable implementations include Cowrie [3], Glutton [4], OpenCanary [7], Conpot [2], T-Pot [10], AIDE [5], ICSNet [19], SIPHON [12] etc. As mentioned, Windows OS are ubiquitous in enterprise IT and (OT environments). Since attacks targeting Windows-based systems have been on the rise in recent years, we provide our review with a focus on related state-of-the-art Windows-based honeypot implementations.

KFSensor, launched in 2003, is one of the first Windows-based honeypot intrusion detection systems to attract and detect hackers and worms by simulating vulnerable system services and trojans [16]. It is available in Professional, Enterprise, and Educational editions. KFSensor could be configured using scenarios, for example, to listen on all TCP and UDP ports, to simulate a MySQL database server or an IIS Web server. Since it simulates vulnerable services, it is likely that attackers will detect that they are interacting with a honeypot rather than a real system. Moreover, KFSensor is a proprietary software that cannot be modified to tailor specific research purposes.

Wang et al. proposed the Strider HoneyMonkey Exploit Detection System, which uses a network of monkey programs running on virtual machines with different patch levels to hunt for websites that exploit browser vulnerabilities [23]. The authors identified 752 unique URLs operated by 287 websites that can successfully exploit unpatched Windows XP machines.

Microsoft has been developing a honeypot sensor network since 2018 [6]. The honeypot framework, written in C#, allows security researchers to quickly deploy various types of exploit handlers, from simple HTTP handlers to complex protocols such as SSH and VNC. In 2021, a dangling subdomain, `code.microsoft.com`, was used temporarily to host a malware C2 service. Instead of removing the subdomain, Microsoft redirected it to the honeypot sensor network till 26 April 2024. The data and findings have been published in [17] and are crucial for understanding the Oday and nDay¹² ecosystem. The framework itself is proprietary, and Microsoft have put in substantial engineering effort into this. Attackers can communicate with over 30 different protocols/services.

HopLab [20], proposed by the Laboratoire de Haute Sécurité (LHS) in Rennes, deploys high-interaction honeypot systems to attract and analyze malicious activities. HopLab is designed for rapid and flexible deployment of honeypots to respond to emerging vulnerabilities such as the Log4j. It can swiftly set up honeypots that emulate these specific weaknesses across various environments, including both Windows and Linux OS enabling timely analysis of new threats as they arise. Although direct comparison with existing state-of-the-art works has been very limiting, **HoneyWin** complements and contributes to ongoing research on high-interaction Windows honeypots, offering an additional layer of monitoring, while also illuminating best practices for honeypot configuration and risk mitigation.

In addition, recent studies further highlight honeypot deployment and deceptive services as means of detecting novel or stealthy attacks. Han et al. [14] present a comprehensive examination of deception techniques within computer security, classifying solutions such as honeypots, honeypots and obfuscation according to various layers (network, system, application, and data) and core objectives (prevention, detection or mitigation of attacks). The survey emphasizes the challenges of generating realistic decoys, optimally placing them, and continuously updating deception elements to avoid attacker evasion. Although the findings show that deception can effectively complement traditional defenses, such as intrusion detection systems, critical gaps remain, including the need for reproducible experiments, robust evaluation methodologies, and reliable ways to seamlessly integrate deception into broader security strategies. **HoneyWin** aims to resolve this issue with a configurable standardized set of deceptive tokens to ensure reproducible results.

7 CONCLUSIONS

We have designed and implemented a high-interaction Windows honeypot system that mimics an enterprise IT environment, namely, **HoneyWin** in this paper. By deploying the **HoneyWin** with three Windows endpoints and an enterprise grade gateway, live in the wild for 34 days it receives more than 5.79 million unsolicited connections, 1.24 million login attempts, 5 and 354 successful logins via RDP and SSH sessions respectively. In addition, the adversary interacted with the deceptive token in one of the RDP sessions and exploited the public-facing Windows endpoint to initiate the SMTP brute-force bot attack. The adversary successfully harvested

1,250 SMTP credentials after attempting 151,179 credentials during the attack. The results indicate that the **HoneyWin** system is enticing for adversaries to probe, compromise, and launch attacks. Moreover, the **HoneyWin** enables us to receive real-time intrusion alerts and attribute attacks via the comprehensive network traffic capturing and host logging capabilities. While we focus the **HoneyWin** deployment in an enterprise environment, we plan to extend this to an OT environment as part of our future work. Ultimately, **HoneyWin** could be scaled and harness the recent development in artificial intelligence technology to realize a robust security posture that detects, analyzes, and mitigates threats effectively.

ACKNOWLEDGEMENTS

This research is supported by the National Research Foundation, Singapore, under its National Satellite of Excellence Programme “Design Science and Technology for Secure Critical Infrastructure: Phase II” (Award No: NRF-NCR25-NSOE05-0001). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

REFERENCES

- [1] [n. d.]. Awesome Honeypots. <https://github.com/paralax/awesome-honeypots>.
- [2] [n. d.]. Conpot - An ICS honeypot. <https://github.com/mushorg/glutton>.
- [3] [n. d.]. Cowrie - An SSH and Telnet Honeypot. <https://github.com/cowrie/cowrie>.
- [4] [n. d.]. Glutton - Generic Low Interaction Honeypot. <https://github.com/mushorg/glutton>.
- [5] Global Cyber Alliance. [n. d.]. AIDE - Automated IoT Defence Ecosystem. <https://globalcyberalliance.org/aide/>.
- [6] Ross Bevington. 2024. Examining the Deception infrastructure in place behind `code.microsoft.com`. <https://techcommunity.microsoft.com/blog/microsoftsentinelblog/examining-the-deception-infrastructure-in-place-behind-code-microsoft-com/4124464>.
- [7] Thinkst Canary. [n. d.]. OpenCanary - A Multi-protocol Network Honeypot. <https://github.com/thinkst/opencanary>.
- [8] Kyle Dickinson. 2020. *Implementer's Guide to Deception Technologies*. Technical Report. <https://www.sans.org/media/analyst-program/implementers-guide-deception-technologies-39390.pdf>
- [9] MITRE Engenuity Center for Threat-Informed Defense. [n. d.]. Defending OT with ATT&CK Reference Architecture. <https://center-for-threat-informed-defense.github.io/defending-ot-with-attack/architecture/>.
- [10] Deutsche Telekom Security GmbH. [n. d.]. T-Pot - The All In One Multi Honeypot Platform. <https://github.com/telekom-security/tpotce>.
- [11] Roger A. Grimes. 2017. *Honeypots*. John Wiley & Sons, Ltd, Chapter 19, 107–110. <https://doi.org/10.1002/9781119396260.ch19>
- [12] Juan David Guarnizo, Amit Tambe, Suman Sankar Bhunia, Martin Ochoa, Nils Ole Tippenhauer, Asaf Shabtai, and Yuval Elovici. 2017. SIPHON: Towards Scalable High-Interaction Physical Honeypots (*CPSS '17*). Association for Computing Machinery, New York, NY, USA, 57–68. <https://doi.org/10.1145/3055186.3055192>
- [13] Fortra Digital Gurdian. 2015. What is the Biggest Misconception Companies Have About Endpoint Security & Protection Tools? <https://www.digitalguardian.com/blog/data-security-experts-answer-what-biggest-misconception-companies-have-about-endpoint-security>.
- [14] Xiao Han, Nizar Kheir, and Davide Balzarotti. 2018. Deception Techniques in Computer Security: A Research Perspective. *ACM Comput. Surv.* 51, 4, Article 80 (July 2018), 36 pages. <https://doi.org/10.1145/3214305>
- [15] Eric Heindl. 2024. The Hidden Dangers of Windows in OT Networks. <https://www.omicroncybersecurity.com/en/resources/the-hidden-dangers-of-windows-in-ot-networks>.
- [16] KeyFocus. [n. d.]. KFSensor: Advanced Windows Honeypot System. <https://www.kfsensor.net/kfsensor/>.
- [17] Francesco Sanna Passino, Anastasia Mantziou, Daniyar Ghani, Philip Thiede, Ross Bevington, and Nicholas A. Heard. 2024. Nested Dirichlet Models for Unsupervised Attack Pattern Detection in Honeypot Data. arXiv:2301.02505 [cs.CR] <https://arxiv.org/abs/2301.02505>
- [18] Fairfield Market Research. 2024. Server Operating System Market. <https://www.fairfieldmarketresearch.com/report/server-operating-system-market>.
- [19] Luis Salazar, Efrén López-Morales, Juan Lozano, Carlos Rubio-Medrano, and Álvaro A. Cárdenas. 2024. ICSNet: A Hybrid-Interaction Honeynet for Industrial

¹²nDay vulnerability is a security flaw that has been disclosed and has an official patch available. However, organizations may still be unpatched and vulnerable.

- Control Systems. In *Proceedings of the Sixth Workshop on CPS&IoT Security and Privacy* (Salt Lake City, UT, USA) (*CPSIoTSec'24*). Association for Computing Machinery, New York, NY, USA, 68–79. <https://doi.org/10.1145/3690134.3694813>
- [20] Alexandre Sanchez. 2024. HopLab: Creating Highly Interactive Honeypots. <https://www.inria.fr/en/hoplab-cybersecurity>.
- [21] SentinelOne. 2024. 7 Types of Ransomware Attacks in 2025. <https://www.sentinelone.com/cybersecurity-101/cybersecurity/types-of-ransomware/>.
- [22] StatCounter. [n. d.]. Desktop Operating System Market Share Worldwide. <https://gs.statcounter.com/os-market-share/desktop/worldwide>.
- [23] Yi-Min Wang. 2005. Strider HoneyMonkeys: Active Client-Side Honeypots for Finding Web Sites That Exploit Browser Vulnerabilities. USENIX Association, Baltimore, MD.
- [24] David Weston. 2024. Helping Our Customers through the CrowdStrike Outage. <https://blogs.microsoft.com/blog/2024/07/20/helping-our-customers-through-the-crowdstrike-outage/>.