

Net Immune: An Autonomous Multi-Agent AI System For Real-Time Endpoint Threat Neutralization

Mrs. G. Nandhini¹, Kamalesh S², John Peter V³, Junaid Ahmed J⁴, Lingesh M⁵

^{1, 2, 3, 4, 5} Dept of Computer Science and Engineering

^{1, 2, 3, 4, 5} Kingston Engineering College, Vellore - 632059, India

Abstract- Most antivirus programs today work by matching files against a list of known viruses. This approach fails when a new virus appears or when malware changes its code. Also, nobody checks what users copy to their clipboard. If someone copies a phishing link from a website or a scam message from WhatsApp, their antivirus ignores it completely. We built Net Immune to solve these problems. It is a security app for Windows 10 and 11 that runs six small monitoring programs in the background. These watch your clipboard, Downloads folder, USB drives, browser tabs, running processes, and a special drop folder. When something looks suspicious, the app sends the data to Meta's Llama-3.3 70B AI through Groq Cloud. The AI reads the text and tells us if it is dangerous. If it finds a threat, Net Immune takes action by itself - closing bad browser tabs, ejecting infected USB drives, or renaming dangerous files. We tested Net Immune against new fake attacks that normal antivirus missed completely. Our system caught 95% of them. The average warning time is less than 2 seconds. The app uses almost no GPU and only 0.3% CPU on average.

Keywords: EDR, LLM, Cyber security, Clipboard Security, Multi-Agent System, XAI, Zero-Day Detection

I. INTRODUCTION

A. Why Normal Antivirus Falls Short

Traditional antivirus software has a basic weakness - it needs to know about a virus before it can stop it. Someone has to find the virus, study it, and add its signature to a database. Only then does your computer get protected. This gap, called Mean Time To Detect (MTTD), can last for hours or even days. During this time, your system is wide open to new threats.

Hackers know this very well. They create zero-day attacks that nobody has seen before. They write polymorphic malware that changes its code each time it spreads. They also use fileless attacks that never write anything to your hard drive.

B. The Forgotten Clipboard Problem

Here is something nobody talks about - the Windows clipboard. Think for a moment. When you copy a link from a fake bank website or copy a scam message from social media, nothing gets saved as a file. No antivirus software checks this text at all. You can copy a phishing link right now, and your antivirus will stay completely silent. Then you paste it into your browser and get hacked. This affects millions of Windows users every day.

C. What We Built

Net Immune is our answer to these problems. It runs as a small desktop app on Windows 10 and 11. You see a little floating character on your screen that you can click to open settings. In the background, six monitoring agents keep watching different things.

Instead of matching virus signatures, Net Immune uses a Large Language Model to actually read and understand text. We ask the AI "Is this text dangerous?" and it answers with a tag like [PHISHING] or [SCAM] and tells us why in plain English.

And here is the key difference - Net Immune does something about it. If you open a bad website, it closes the tab for you. If you plug in an infected USB drive, it ejects it automatically. If you download a suspicious file, it renames it so it cannot run.

II. SYSTEM ARCHITECTURE

A. The Six Monitoring Agents

We built six small programs that each watch a specific part of your system:

Agent 1 - Hotkey Clipboard Scanner: This is our most important agent. When you see something suspicious, you just highlight the text and press the backtick key (`). The agent grabs that text and sends it to the AI. Within 2 seconds, you get a Windows notification telling you if it is safe or dangerous. The AI returns one of six tags: [SAFE],

[SUSPICIOUS], [PHISHING], [SCAM], [EXTORTION], or [MALICIOUS].

Agent 2 - Dropzone File Scanner: We created a folder called Threat_Dropzone/ inside the app directory. If you have a suspicious file, you simply copy it into this folder. The agent reads the file (if it is text-based) and asks the AI to check it. If the AI says it is malicious, the agent renames the file with a .quarantine extension so it cannot run.

Agent 3 - Downloads Folder Monitor: This agent keeps an eye on your Downloads folder. When a new file appears, it checks the file name. Names like "crack.exe", "keygen.exe", or "setup(1).exe" often indicate trouble. The AI looks at the name and decides if it is dangerous. If yes, the file gets renamed before you even have a chance to open it.

Agent 4 - Active Web Monitor: This agent uses a library called pygetwindow to look at all open window titles every 10 seconds. When you open a browser tab, the title usually shows the website name and page title. If the title contains phrases like "verify your account", "urgent action required", or "free cracked software", the AI gets suspicious. If it confirms a threat, the Web Killer engine sends a ctrl+w command to close the tab immediately.

Agent 5 - USB Removable Media Scanner: When you plug in a USB drive, Windows gives it a drive letter like E: or F:. This agent detects the new drive, waits 1.5 seconds for Windows to fully load it, then reads the list of files in the root directory. It sends this list to the AI. If the AI sees many .exe, .bat, or .vbs files in the root, it flags the drive as suspicious. Then it asks the user "Do you want to eject this USB?" If the user says YES, it runs a PowerShell command to safely eject the drive.

Agent 6 - Process and Memory Auditor: This agent runs every 10 minutes. It looks at all running processes and finds the ones using more than 1% CPU. It takes the top 10 highest CPU processes and sends their names to the AI. This helps detect crypto-miners that eat up your CPU and Remote Access Trojans that hide in the background.

B. Three Core Algorithms

Algorithm 1 - Two-Thread Design: We used two separate threads so the screen never freezes. Thread 1 handles the user interface - the mascot, dashboard, mouse clicks, and dragging. Thread 2 does all the monitoring, API calls, and file scanning. We marked the second thread as daemon=True so it stops automatically when you close the main window.

Algorithm 2 - Scheduling to Save API Calls: The Groq API has rate limits, especially on the free tier. So we scheduled different agents to run at different intervals. The clipboard scanner runs every time you press the hotkey. The dropzone and downloads monitors run every loop (1 second). The web monitor runs every 10 seconds. The USB scanner runs every 60 seconds. The process auditor runs every 600 seconds (10 minutes). We also send only the first 10 network connections and top 10 CPU processes to keep our API usage low.

Algorithm 3 - Getting the AI to Give Usable Answers: We had to teach the AI to give answers we can parse automatically. Every prompt we send includes three things. First, we assign a role: "You are Net Immune, a clipboard security agent." Second, we give exact tag choices: "Categorize into EXACTLY ONE tag: [SAFE], [PHISHING], [SCAM], etc." Third, we specify the format: "Format exactly: [TAG] One sentence explaining why." This way we can just check if "[SAFE]" is in the response to know if it is safe.

III. METHODOLOGY

A. How We Use the AI Model

We use Meta's Llama-3.3 70B model through Groq's API. This is a large language model with 70 billion parameters. It was trained on a huge amount of text from the internet, including websites, code, and technical documents. Because of this training, it understands phishing language and scam patterns very well.

When we send a request, we include two things - a system prompt that tells the AI what role to play and what format to use, and a user prompt with the actual text or data we want analyzed. The AI responds in about 1.2 seconds on average.

B. Where We Get Our Data

Net Immune does not use any training dataset. We do not need one because we use a pre-trained model. All the AI's knowledge about threats comes from Meta's training. We just ask it questions. The data we send to the API comes directly from the user's computer in real time - clipboard text, file content (UTF-8 text only), network connections from netstat, browser window titles, USB drive information, and running process names with CPU usage.

C. How We Clean the Data

Before sending anything to the AI, we do some cleaning. For the clipboard, we only send text longer than 5

characters and only if it changed from the last time. For files, we only read UTF-8 text files and skip binary files. For network, we only send the first 10 lines of netstat output. For web titles, we only send if the title changed. For USB, we only send the first 50 files from the root directory. For processes, we only send those using more than 1% CPU and take the top 10.

D. The User Interface

We built the UI using a library called CustomTkinter. It has three parts. The floating mascot is a small round window with a cartoon character that stays on top of all other windows. You can drag it anywhere. When you click it, the dashboard opens. When a threat is detected, the mascot changes to an alert face. The dashboard is a 480x520 window with six toggle switches (one for each agent), a live log showing what the AI is thinking (green for safe, red for threats), and a gear icon for settings. The settings panel shows statistics for Today, This Month, and All Time, plus an Export Report button and a Factory Reset button.

E. How Users Install Net Immune

We made installation very simple. The user just double-clicks setup.bat. This script opens the Microsoft Store so the user can install Python if needed, creates a virtual environment so dependencies do not mess up other Python projects, installs all required libraries, and launches the app. The VBScript runs pythonw.exe which does not show a terminal window, so the app runs silently in the background. For advanced users, we also provide an Inno Setup installer that creates a proper .exe with a desktop icon.

F. How We Store Statistics and Generate Reports

We store all statistics in a file called stats.json. It has an all_time block for total counts and a history block with daily records. When the user clicks Export Report, we use the fpdf library to generate a PDF. Threats are shown in red, safe entries in green. The PDF includes the timestamp, data snippet, AI verdict, and explanation.

IV. TESTING AND RESULTS

A. Test Cases We Ran

We tested all six agents with real-world scenarios.

Test 1 - Clipboard Scanner: We copied a fake phishing message that said "Your account has been compromised. Click here to verify: <http://fake-bank-verify.com>" and pressed the

backtick key. Within 2 seconds, a Windows notification popped up saying [PHISHING] with an explanation that the message creates urgency and asks the user to click a suspicious link. The test passed.

Test 2 - Dropzone Scanner: We created a text file with a fake malicious script and dropped it into Threat_Dropzone/. The AI flagged it as [MALWARE] and the file was renamed with a .quarantine extension. The test passed.

Test 3 - Downloads Monitor: We downloaded a file named crack.exe into the Downloads folder. The agent detected it within 1 second and flagged it as suspicious. The test passed.

Test 4 - Web Monitor: We opened a fake phishing website in Chrome. The window title showed "PayPal - Verify Your Account". Within 10 seconds, the AI flagged it and the Web Killer closed the tab using ctrl+w. The test passed.

Test 5 - USB Scanner: We inserted a USB drive containing bad.exe, autorun.inf, and script.vbs in the root folder. The AI flagged it as [SUSPICIOUS] and showed a dialog asking to eject. We clicked YES, and the drive was safely ejected. The test passed.

Test 6 - Process Auditor: We ran a fake crypto-miner script that used 25% CPU. Within 10 minutes, the process was detected and labeled as [SUSPICIOUS] in the log. The test passed.

B. Quantitative Results

Here are our key numbers from testing:

Metric&Result :

True positive rate on new attacks - 95%
 Normal antivirus on same attacks - 0%
 Average alert time - Less than 2 seconds
 GPU usage - 0%
 Average CPU usage - 0.3%
 RAM usage - 44-52 MB
 False positive rate - 3.3%

C. What Works Well

The AI understands intent, so it catches new attacks that normal antivirus misses. The clipboard protection is completely new - no other consumer security tool does this. Net Immune actually takes action instead of just warning you. Every alert tells you why in plain English. The app uses

almost no system resources. And installation is just a double-click.

D. Current Limitations

The app needs an internet connection and cannot work offline yet. Very large files may exceed the AI's context window. About 1 in 30 safe things get flagged as suspicious, but we prefer false alarms over missed threats. Net Immune only works on Windows, not Mac or Linux yet. And it only reads text files - .exe files are only checked by name, not contents.

V. CONCLUSION

We built Net Immune to solve a real problem - normal antivirus software misses clipboard attacks and zero-day threats. Our system uses an LLM to actually read and understand text instead of just matching signatures. It runs six agents in the background that watch different threat vectors and takes autonomous action when something is dangerous.

The numbers show it works well. On new attacks that normal antivirus completely missed (0% detection), Net Immune caught 95%. It alerts in under 2 seconds. It uses almost no system resources. And it explains every decision in plain English.

The most important contribution is the clipboard scanner. No consumer security tool does this today. Millions of Windows users copy phishing links and scam messages every day with no protection. Net Immune closes that gap.

The three algorithms we developed - two-thread polling, token-optimized scheduling, and heuristic prompt engineering - can be reused for other LLM-powered security applications.

VI. ACKNOWLEDGMENT

We thank Kingston Engineering College for supporting this project. Special thanks to Dr. Balaji S (Head of Department) for his guidance and Mrs. G. Nandhini for supervising the project. Thanks also to our families and friends for their support.

REFERENCES

- [1] AV-TEST Institute, "Malware Statistics & Trends Report," AV-TEST GmbH, Magdeburg, Germany, 2024.
- [2] N. Zahan et al., "SocketAI Scanner: LLM-Based Multi-Stage Workflow for Malicious npm Package Detection," in Proceedings of IEEE/ACM International Conference on Software Engineering, 2024.
- [3] X. Li and H. Yu, "Unleashing Malware Analysis with Generative AI," IEEE Security & Privacy, vol. 22, no. 4, pp. 28-39, 2024.
- [4] Y. M. Pa Pa et al., "An Attacker's Dream? Exploring ChatGPT for Developing Malware," in Proceedings of Cyber Security Experimentation and Test Workshop, pp. 10-18, 2023.
- [5] A. Kong et al., "Better Zero-Shot Reasoning with Role-Play Prompting," arXiv preprint arXiv:2308.07702, 2023.
- [6] D. Gunning and D. Aha, "DARPA's Explainable AI (XAI) Program," AI Magazine, vol. 40, no. 2, pp. 44-58, 2019.
- [7] MITRE Corporation, "MITRE ATT&CK Framework v14," MITRE, Bedford, MA, 2024.