

Malware Analysis: Quantifying Attack Impact on System Resources

Senior Thesis

By: Fallyn Buckner

Submitted to: Mark Huber

This research analyzes the potential for system performance metrics to function as indicators of compromise for Zeus malware infection. The research question addressed is: “To what degree can changes in CPU, RAM, and network metrics be used to detect Zeus malware with statistical significance?” Due to institutional approval constraints, this study presents a methodological framework and projected outcomes based on existing literature rather than executed experiments. This study will examine clear measurable changes in CPU, memory consumption, and network activity before and after infection in a sandbox environment. The objective is to determine if general system monitoring is sufficient to serve as a security layer in malware detection programs by analyzing the statistical significance of observed changes in system performance.

December 1, 2025

Contents

1	Introduction	3
2	Literature Review	4
2.1	System Performance Indicators for Malware Detection	5
2.2	Detection Methodologies and Their Effectiveness	6
2.2.1	Optimizing Monitoring Environments and Analysis Tools	7
2.3	Limitations of Current Approaches	8
2.4	Relationship to Current Research	9
3	Methodology	10
3.1	Preliminary Approvals	10
3.2	Incident Response Plan	10
3.3	Containment Procedures	13
4	Separation Validation Summary	13
5	Experimental Setup	14
5.1	Malware Analysis Security Architecture	14
5.2	Malware Sample	15
5.2.1	Isolation Architecture	15
5.3	Hardware Configuration	15
5.3.1	Physical Layer	15
5.3.2	Cloud Compute Resources	16
5.3.3	Analysis Target (Windows Server 2019)	17
5.3.4	Monitoring Tools	17
5.4	Metrics Collection	17
5.5	Measurement Tools	18

5.6	Experimental Protocol	19
6	Projected Results Based on Literature	21
6.1	Preliminary Findings	21
7	Data Analysis Plan	24
7.1	Data Preprocessing	24
7.2	Statistical Tests	24
7.3	Visualization Approaches	25
7.3.1	Primary Visualization: Bar Charts with Confidence Intervals	25
7.3.2	Supplementary Visualization: Time Series Analysis	27
8	Study Limitations	28
9	Conclusion	28

1. Introduction

Zeus is a type of malware that has been historically used to steal banking credentials. Specifically, it is categorized as a trojan. It is named after the famous Greek trope. The malware enters a system packaged as non-malicious software designed to accomplish a useful task, but the internal contents of the program have instructions to carry out malicious behavior under the scenes. The research presented in this project will rely mainly on dynamic malware analysis. With the research question, “To what degree can changes in CPU, RAM, and network metrics be used to detect Zeus malware with statistical significance?” in mind, the objective of analyzing current research is to determine if general system monitoring is sufficient for malware detection programs.

Cybersecurity is an expansive field that has proven its need in society. This has especially held true in recent years. There is not a month that goes by where news outlets fail to mention account compromises, data leaks, and instances of fraud. All these travesties were accomplished through technology. The benefits and conveniences of technology have attracted individuals wishing to take advantage of innocent people. Individuals not educated or prepared enough to handle the challenge of securing their digital assets fall victim the most. Pairing convenience with security will be one of the toughest challenges of this digital age.

The security landscape is a playing field where companies must cover all bases: network access, account protection, social engineering. Adversaries must only find one vulnerability to penetrate company infrastructure. The balance between the two groups is not exactly fair. Hiring knowledgeable and forwardthinking security staff becomes a nonnegotiable for companies who wish to keep the trust of consumers. Cyberattacks tarnish the reputation of organizations leading consumers to think twice about keeping their personal information within that organization. Compliance is the minimal level a company must reach to operate in the eyes of the government, but reaching this level will not fully protect against the

imminent threat of cyberattacks.

Security Engineers become the guardians of consumer data. Quality ones will anticipate potential threats and formulate systems for mitigating given risk in an organization. Organizations hit with malware or ransomware related cyberattacks must be able to reverse engineer the malware to figure out how it was able to infiltrate systems and if any data has been exfiltrated.

Security research ultimately aids in creating more proactive systems that can protect against threats based on existing knowledge. The information about how malware will behave provides something known as a signature. Behaviors are things such as changing a specific file name, making a C2C call to a given IP, or moving a registry to an alternative location. How a malware program behaves distinguishes it among other types of malware. Security programs like antivirus can proactively identify and remediate the malware because the software knows exactly what to look for. Training for such a purpose was done in the hands of security researchers.

2. Literature Review

Existing research and scholarly work provide a framework for how this research should be conducted. Three key works have integrated important perspectives to analysis and planning for research on malware detection through system performance metrics. Kaushik et al. (2024) in “Advanced Techniques and Applications of Cybersecurity and Forensics” provided general baseline information on malware types and Zeus-specific system affecting metrics. Basu et al. (2019) with “A Theoretical Study of Hardware Performance Counters-Based Malware Detection” added claims about ensuring sampling standards provide valuable information and how multiple indicator metrics will provide best results. Additionally, Luo et al. (2022) in “On Teaching Malware Analysis on Latest Windows” highlighted the importance of testing on Windows devices and offered a practical framework for actual analysis in sandbox environments. The culmination of these three works will ultimately steer the

direction of providing analysis where clear conclusions can be drawn.

2.1. System Performance Indicators for Malware Detection

A clear consensus of system metrics being essential to coming up with a baseline reading of malware presence is important was present in all works. Computer metrics like CPU, RAM, and network activity indicate what programs are running and how much of the systems resources are being allocated to them. If harmful programs are present on the system, quantitative indicators of such will be present in these metrics. That is what makes them invaluable.

The objective of a specific malware type may affect how the malware presents itself through its interactions on a host device. Coupling this with the fact that according to Kaushik et al. (2024, p. 174), the computing power of it is exponentiated through the use of a botnet, the scale of Zeus even more alarming. Kaushik et al. identifies that Zeus “tries to steal credentials online by capturing keystrokes and screenshots, injecting HTML into web pages, and exploiting vulnerabilities present in a browser. The server sends a file that consists of lists of actions to be performed on the victim’s machines to the bot so that it can do all the activities according to the file.”

Zeus’ objective is data theft through various attack vectors. It accomplishes these objectives through keylogging for credential harvesting and numerous network-based exploitation techniques, and “injecting HTML into web pages, and exploiting vulnerabilities present in a browser.” Kaushik et al. (2024, p. 174) also notes that, some variants encrypt the data and act as ransomware, depending on configuration. The authors also outline the exact system impacts that one can expect if their device is infected by Zeus. Reports find that there are impacts in the user system experience such as degradation in the speed of the operating system, lagging while doing tasks, and the presence of suspicious processes indicating command and control communications occurring in the background. These traits differ from the system resources category because they are what the end user experiences. The system itself

is noted to have “consumption of CPU power; the heating of the system itself presented in hardware is” a major red flag” according to Kaushik et al. (2024, p. 172). All these technological impairments are noticeable, but the effects in the user are shown in their account. Unusual transactions in bank accounts are a common indicator of compromise (IOC) that prove successful credential theft and financial fraud perpetrated by the threat actor who leveraged the Zeus banking Trojan. The financial nature of target components is shown in this instance

2.2. Detection Methodologies and Their Effectiveness

There are two specific types of analysis when it comes to malware. Quality analysis consists of both. Dynamic and static analysis bring benefits that have aided tremendously in understanding malware operations on a granular level. Static analysis consists of reviewing the codebase of a malware program without running it. It examines the properties of what the malware is programmed to do. Static analysis has significant limitations because sophisticated malware often uses “various techniques (obfuscation) to hide what it is actually doing” (Kaushik et al., 2024, p. 3). Static analysis can provide only the logic behind the layers of confusion that threat actors seek to produce. Integrating these findings can help inform dynamic analysis. It is an important step in the overall research process because malware has an important nuance. Not all malware presents itself uniformly on all systems. Some malware changes its course of action based on host system constraints and programming by threat actors. Therefore, static analysis is the only form of analysis that can provide a full outlook on the range of what the malware has the capacity to accomplish.

However, its counterpart, dynamic analysis, has key benefits in other areas. These benefits make it a more impactful type of research in the case of Zeus. Luo et al. (2022, p.4) explain the necessity of dynamic techniques: “basic static techniques really only provide some hints about the executable,” making runtime observation essential. This approach is grounded in observed behavior. It heavily supports the protocol of monitoring system met-

rics during execution to detect anomalous behavior. Hints provided in static analysis can be fully actualized in dynamic analysis. Researchers can see impacts just as infected users would in the wild.

2.2.1 Optimizing Monitoring Environments and Analysis Tools

For malware analysis, Luo et al. (2022, p. 2) emphasize the necessity of “a virtual machine (VM) environment” where one “can roll back to a previous good snapshot” when infected. The controlled sandbox approach will ensure safe performance monitoring without jeopardizing necessary production systems. There are foundational types of monitoring tools that researchers find to be essential to the process. Process monitors, process explorers, registry analysis tools, and network analysis tools make the core team.

Their contributions are as follows:

1. Process Monitors provide specific information like timestamps, processes being ran, and specifics about operations performed.
2. Process explorers give information about running processes in a tree-like hierarchy and show relationships between processes. The tree-like visual element can aid in mapping how events correlate to one another.
3. Registry analysis tools aid in comparative efforts by showing snapshots before and after execution. Having clear metrics of system performance metrics before pivotal changes can be especially useful in analysis efforts. The registry analysis streamlines processes.
4. Network analyzers are interactive interfaces that provide a report of network traffic and destinations. Some of the key metrics include IP addresses, domain names, port numbers, MAC addresses, host names, and connections.

2.3. Limitations of Current Approaches

There are clear limitations on how existing research can be used for current operations. Hardware Performance Counters (HPC) based malware detection research has mainly been focused on rootkits rather than banking trojans. As Basu et al. note, their methodology was validated against rootkits that “modify system libraries or kernel-level behavior” (2020, p. 514). Zeus malware operates primarily as a banking trojan with different behavioral characteristics, as “not all variants necessarily access privileges up to administrator level” (Kaushik et al., 2024, p. 172). This presents a limitation on how the research methodologies can be chained together. These discrepancies cause tension in detection. Banking trojans like Zeus focus on data theft through keylogging and form grabbing rather than deep system manipulation and potentially create different performance signatures than the rootkits studied in existing, referenced literature.

Another limitation in research is how signature-based detections are integrated. These types of systems demonstrate significant weaknesses across different Zeus variants. Basu et al. emphasize that “although anti-virus software is popular, an ongoing cat-and-mouse cycle of anti-virus development and malware that thwarts the anti-virus has ensued” (2020, p. 512). This limitation stems from the fundamental vulnerability that “software can be tricked and evaded” which is exactly why “finding hardware-based detection models and having them built in is important” (Basu et al., 2020, p. 512). In this area, more research needs to be done to see how Zeus performs and interacts with hardware-based analysis measures. Intelligent adversaries can develop evasion techniques against signature-based detection with ease. The ways that researchers can receive different responses to system manipulations from virus infections reflect this. Luo et al. has evidence that takes a similar stance. Luo et al notes that researchers should “disable all Windows 10 exploit protection settings, turn all Windows Defender Firewall settings on and turn all the Virus threat protection settings on” (2022, p. 3) during their analysis work. These limitations collectively point to the need for multi-faceted approaches that combine system metric monitoring with traditional detection

methods. With both facets integrated, a more secure defense against banking trojans like Zeus can start to take form.

2.4. Relationship to Current Research

Current research provides insight into the specific considerations. The consistent standard of analysis that the papers follow show coherency in how malware detection is looked upon. Having a statistical analysis framework for comparison is important, so choosing variables such as “degradation in the speed of the operating system” and “consumption of CPU power” (Kaushik et al., 2024, p. 172) in consistent ratings like “comparing the registry by taking two snapshots, one before and one after running an executable” (Kaushik et al., 2024, p. 187) pave the way for clear conclusions to be made. This research goes further by specifically focusing on how “dynamic analysis can be divided into two parts host-based and network-based indicators” (Luo et al., 2022, p. 2-3).

The topic itself is unique because it frames the malware in the context of impact, directly bridging and quantifying how its results manifested in the financial industry. While existing research has established detection frameworks and identified system metrics affected by malware, there remains a significant gap in quantitatively measuring the statistical significance of these system performance changes specifically for Zeus Trojan detection. Previous studies have focused primarily on rootkits rather than banking trojans (Basu et al., 2020) or have provided general detection frameworks without establishing precise statistical thresholds for system metric deviations that indicate Zeus infection. This research extends the findings of Kaushik et al. (2024) and Luo et al. (2022) by specifically quantifying the degree to which CPU, RAM, and network metric changes can reliably indicate Zeus infection. It could potentially enable more accurate automated detection systems. The integration of multiple system indicators analyzed together, rather than looking at single metrics, creates a stronger framework for detection that can lower false positives while still catching real threats. This approach gives financial systems better protection against Zeus and similar advanced Tro-

jans that continue to evolve in our digital landscape. Learning from the past is the best way to ensure better efforts for protecting the future.

3. Methodology

This research employs a controlled experimental approach to measure and analyze system performance metrics before and after malware infection.

3.1. Preliminary Approvals

This research was designed as an expository study due to institutional approval constraints. The experimental framework presented here follows AWS Security Blog standards for isolated malware research and demonstrates the theoretical approach that would be required for actual malware analysis in an academic setting.

In a full implementation, the following approvals would be required:

- CMC Information Security Department architecture review and written approval for malware research
- AWS account authorization for dedicated security research environment
- End-of-lifecycle device allocation for additional isolation from institutional systems

The methodologies, security controls, and statistical analysis frameworks presented throughout this paper represent best practices for this type of research and could be directly implemented given appropriate institutional clearance.

3.2. Incident Response Plan

A remediation plan is the safety net for possible misdrifts or improper functionality. A plan provides researchers with an action list and point of reference should analysis go wrong at a point. To best streamline the process for researchers having a clear, consistent

framework to begin with, security hardening is key. Security hardening is the process of putting in security protocols for a system such as updates and configuration preferences. A daily verification script to automatically check that the environment within the VPC is hardened and restricted as expected would be ideal to accomplish the desired task. If the preliminary criteria of the script is not met, a warning will be presented to the researcher.

Daily Verification Script (run before each experiment):

```
#!/bin/bash
# Verify network isolation

# Check for Internet Gateway
IGW=$(aws ec2 describe-internet-gateways
  --filters "Name=attachment.vpc-id,Values=$VPC_ID"
  --query 'InternetGateways[0].InternetGatewayId'
  --output text)
if [ "$IGW" != "None" ]; then
  echo "ERROR: Internet Gateway detected! Experiment aborted."
  exit 1
fi

# Check for NAT Gateway
NAT=$(aws ec2 describe-nat-gateways
  --filter "Name=vpc-id,Values=$VPC_ID" "Name=state,Values=available"
  --query 'NatGateways[0].NatGatewayId'
  --output text)
if [ "$NAT" != "None" ]; then
  echo "ERROR: NAT Gateway detected! Experiment aborted."
  exit 1
fi

# Check for default route
ROUTE=$(aws ec2 describe-route-tables
  --filters "Name=vpc-id,Values=$VPC_ID"
  --query 'RouteTables[*].Routes[?DestinationCidrBlock=='0.0.0.0/0']'
  --output text)
if [ -n "$ROUTE" ]; then
  echo "ERROR: Default route (0.0.0.0/0) detected! Experiment
    aborted."
  exit 1
fi
```

```
echo "Network isolation verified. It is safe to proceed with testing  
..."
```

3.3. Containment Procedures

Critical Incident (Evidence of campus network exposure):

```
Immediate Actions (<5 minutes):
1. Terminate all EC2 instances via AWS Console
2. Verify VPC route tables (confirm no IGW exists)
3. Review VPC Flow Logs for any unexpected ACCEPT actions
4. Check CloudTrail for unauthorized API calls (VPN creation,
   peering, etc.)
5. Contact institutional CISO: [phone/email]
6. Preserve logs (CloudTrail, VPC Flow Logs, GuardDuty findings)
```

Forensic Analysis:

- Export CloudTrail logs for past 90 days
- Analyze VPC Flow Logs for any traffic to campus IP ranges
- Review IAM Access Analyzer for cross-account access patterns
- Generate timeline of all EC2/VPC API calls

4. Separation Validation Summary

```
Layer 1: ACCOUNT SEPARATION
        Dedicated AWS account
        No organizational enrollment
        Separate billing

Layer 2: NETWORK SEPARATION
        No internet gateway
        No NAT gateway
        No VPN/peering connections
        DNS sinkhole (non-routable responses)
        Security groups (deny-all egress)

Layer 3: ACCESS SEPARATION
        Session Manager only (no SSH/RDP)
        IMDSv2 with hop limit = 1
        Permission boundaries (no IAM escalation)
        Ephemeral instances (single-use)
```

Malware Analysis Security Architecture

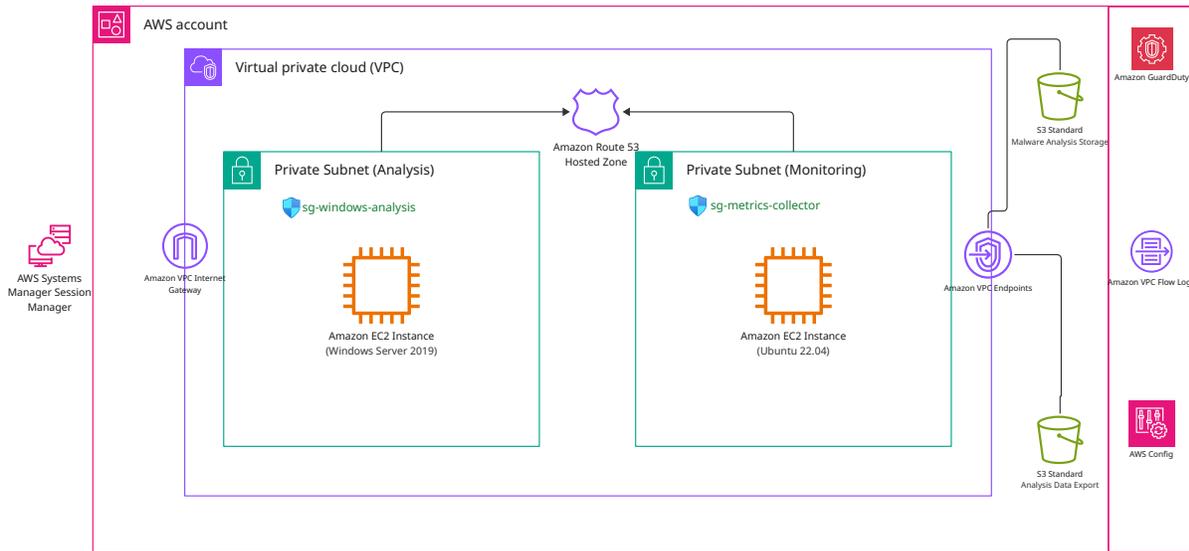


Figure 1: AWS Malware Analysis: Security Architecture

5. Experimental Setup

5.1. Malware Analysis Security Architecture

This research proposes a cloud-based malware analysis environment following AWS Security Blog standards for isolated malware research. The architecture employs six layers of isolation to ensure zero connectivity between the malware analysis environment and institutional networks. The complete architecture is illustrated in Figure

Configuration Example

```
VPC Flow Logs
  Capture: ALL traffic (accepted + rejected)
  Destination: CloudWatch Logs
  Format: Default (14 fields including 5-tuple, bytes, packets)
  Aggregation: 60 second intervals
  Retention: 7 days
```

Possible Findings Example

```
GuardDuty Findings (Expected Baseline):  
  Backdoor:EC2/C&CActivity.B - Zeus C2 attempts (EXPECTED)  
  CryptoCurrency:EC2/BitcoinTool.B - Potential cryptomining (  
  EXPECTED)  
  UnauthorizedAccess:EC2/TorRelay - Tor network usage (POSSIBLE  
  )
```

5.2. Malware Sample

Zeus malware sourced from VX underground will be used for this study. VX Underground is a research organization and online repository that archives malware samples and other information about malicious software for educational purposes.

5.2.1 Isolation Architecture

The architecture would implement six layers of defense isolation as summarized in Table 1.

5.3. Hardware Configuration

5.3.1 Physical Layer

Malware analysis conducted on dedicated end-of-lifecycle laptop for additional isolation from institutional systems.

Isolation Layer	Key Controls
Account Separation	<ul style="list-style-type: none"> • Dedicated AWS account • Independent authentication and billing • No organizational enrollment
Network Separation	<ul style="list-style-type: none"> • No IGW/NAT Gateway/VPN • Local routes only • Deny-all security groups and NACLs • DNS sinkhole (Route53 Private Hosted Zone)
Compute Separation	<ul style="list-style-type: none"> • IMDSv2 enforced (hop limit = 1) • Ephemeral instances (24 hour max lifetime) • Complete termination after each session
Data Separation	<ul style="list-style-type: none"> • Isolated S3 buckets (read-only samples, write-only data) • AES-256 encryption • VPC Gateway Endpoint for private access
Identity Separation	<ul style="list-style-type: none"> • IAM permission boundaries prevent escalation • Limited to CloudWatch, S3, logs • Explicit denial of IAM/network modifications
Monitoring Controls	<ul style="list-style-type: none"> • AWS Config (drift detection) • VPC Flow Logs (60-sec intervals) • GuardDuty (threat detection)

Table 1: Six-Layer Isolation Architecture

5.3.2 Cloud Compute Resources

Implementation:

```

Analysis Session Lifecycle:
1. Launch EC2 from clean AMI baseline
2. Install monitoring agents (CloudWatch, Sysmon)
3. Take EBS (Elastic Block Store) snapshot to be pre-infection
   baseline
4. Execute Zeus malware sample
5. Collect metrics for 4 hours
6. Export data to S3
7. TERMINATE instance (not stop - full destruction)
8. Delete EBS volumes
9. Repeat for next analysis iteration

```

Specification	Analysis Target	Monitoring Instance
Platform	Windows Server 2019	Ubuntu 22.04 LTS
Instance Type	t3.medium	t3.small
vCPUs	2	2
Memory	4 GiB	2 GiB
Storage	30 GiB (gp3 EBS)	20 GiB (gp3 EBS)
Network	Up to 5 Gbps	Up to 5 Gbps
Purpose	Malware infection target	Metrics aggregation

Table 2: EC2 Instance Specifications

5.3.3 Analysis Target (Windows Server 2019)

Zeus targets Windows systems. Windows Server 2019 represents enterprise environments typically affected by banking trojans.

Configuration:

- Windows Defender disabled to prevent malware quarantine
- Windows Firewall configured for monitoring traffic only
- Standalone workgroup so there is no domain join
- No additional software beyond monitoring tools

5.3.4 Monitoring Tools

Tool	Version/Interval	Data Captured
Sysmon	v14.16	Process creation, network connections, file creation, registry modifications
CloudWatch Agent	60 seconds	CPU utilization, memory utilization, disk I/O, network I/O
Performance Monitor	1 second	Processor time, available memory, network bytes/sec, process handles

Table 3: Installed Monitoring Tools

5.4. Metrics Collection

The following metrics will be collected before and after infection:

1. CPU usage percentage
2. RAM/Memory consumption
3. Network metrics:
 - a. Bandwidth usage
 - b. Packet rate (packets per second)
 - c. Connection count (number of active connections)
 - d. DNS query frequency

5.5. Measurement Tools

Metric	Tool/Source	Interval	Measurement
3*CPU Utilization	CloudWatch	60 sec	% utilization (0-100%)
	Performance Monitor	1 sec	Processor time counter
	Sysmon Event ID 1	Real-time	Per-process CPU time
2*Memory Utilization	CloudWatch	60 sec	Total, available, used memory
	Performance Monitor	1 sec	Working Set, Private Bytes, Virtual Bytes
4*Network Activity	VPC Flow Logs	60 sec	Connection count (unique 5-tuples/min)
	VPC Flow Logs	60 sec	Bandwidth (bytes/min → MB/s)
	VPC Flow Logs	60 sec	Packet rate (packets/sec)
	VPC Flow Logs	60 sec	C2 attempts (REJECT actions/min)
	Sysmon Event ID 3	Real-time	Network connections with process correlation
	Route53 Query Logs	Real-time	DNS query frequency

Table 4: Performance Metrics and Measurement Tools

5.6. Experimental Protocol

Each experimental session consists of two hour observation periods: baseline (pre-infection) and infected (post-infection). The complete protocol is summarized in Table 5.

Phase	Duration	Key Activities
Setup		
Pre-Experiment Validation	15 min	Verify isolation (no IGW/NAT), AWS Config compliance, GuardDuty active, S3/CloudWatch operational
Instance Provisioning	10 min	Launch Windows & Ubuntu instances; verify monitoring tools; create EBS snapshot
Baseline Phase		
System Stabilization	30 min	Windows boot processes complete; system reaches idle state (<10% CPU)
Baseline Measurement	3.5 hours	Collect metrics with no user interaction; validate continuous data collection
Data Export	5 min	Export all metrics to S3 /baseline/ directory
Infection Phase		
Pre-Infection Snapshot	5 min	Final EBS snapshot; export process list, registry baseline, network connections
Zeus Detonation	10 min	Download sample from S3; verify hash; execute as admin; document timestamp
Infection Verification	30 min	Confirm Zeus process active; C2 attempts in VPC Flow Logs; GuardDuty findings
Infected Measurement	4 hours	Collect metrics during infected state; hourly validation of Zeus activity
Data Export	5 min	Export all metrics to S3 /infected/ directory
Cleanup		
Instance Termination	5 min	Terminate both instances; delete EBS volumes; verify deletion
Data Preprocessing	30 min	Align time series; calculate derived metrics; standardize format
Total Session Duration: ~9 hours		

Table 5: Experimental Protocol Summary

Malware Samples Bucket:

```
Bucket: malware-samples-bucket-<account-id>  
Encryption: AES-256 (AWS-managed keys)  
Versioning: ENABLED  
Public Access: BLOCKED (all 4 settings)  
Bucket Policy: Deny non-HTTPS uploads  
Lifecycle: Delete after 90 days (samples destroyed)
```

Analysis Data Bucket:

```
Bucket: analysis-data-bucket-<account-id>
Encryption: AES-256 (AWS-managed keys)
Versioning: ENABLED
Public Access: BLOCKED
Contents: CloudWatch metrics, VPC Flow Logs, Sysmon events
Lifecycle: Transition to Glacier after 90 days
```

Complete protocol repeated for five sessions consisting of five baseline and five infected observations) with 24 hour cooling periods between sessions for data quality review and isolation control verification.

6. Projected Results Based on Literature

6.1. Preliminary Findings

Based on Zeus malware characteristics documented in the literature review and established behavioral patterns of banking trojans, the following performance metric changes are anticipated during controlled infection:

Metric	Baseline (Pre)	Infected (Post)	Change	Contributing Factors
CPU Utilization	5-15%	25-45%	+20-30pp	<ul style="list-style-type: none"> • Keylogging • Form grabbing • C2 comms
Memory	1.5-2.0 GB	2.5-3.5 GB	+1.0-1.5 GB	<ul style="list-style-type: none"> • Zeus process • Browser injection • Credential storage
Connections	Baseline	3-5x baseline	3-5x	<ul style="list-style-type: none"> • C2 beaconing • Exfiltration • DGA activity
DNS Queries	Baseline	2-4x baseline	2-4x	<ul style="list-style-type: none"> • DGA for C2
Packet Rate	Normal	Sustained elevation	Continuous	<ul style="list-style-type: none"> • C2 comms • Credential transmission
Bandwidth	Baseline	Periodic spikes	Burst patterns	<ul style="list-style-type: none"> • Screenshots • Form data • Keystroke logs

Table 6: Expected Performance Metrics: Baseline vs. Infected State

Metric	Why This Change Is Expected
CPU Utilization	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Continuous keyboard monitoring for keylogging • Real-time browser manipulation for form grabbing • Encryption/decryption for C2 communication <p>As Kaushik et al. (2024, p. 172) note, “consumption of CPU power” is a documented indicator of Zeus infection. The malware’s multi-threaded design distributes these tasks across available CPU resources, resulting in measurable increases.</p>
Memory Consumption	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Core malware process (50-150 MB) • Code injected into browser processes (200-400 MB per active browser) • Credential storage buffers • Screenshot caching <p>The malware injects itself into legitimate processes to maintain persistence. The memory footprint of infected applications is impacted as a result.</p>
Connection Count	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Persistent connections for command polling (every 2-5 minutes) • New connections for credential exfiltration • Reconnaissance traffic to identify banking domains <p>Baseline idle systems typically have 10 to 20 established connections, which Zeus multiplies through its distributed C2 architecture.</p>
DNS Query Rate	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Domain Generation Algorithms (DGA) to locate active C2 servers • 50-200 domain queries per hour generated by DGA <p>Baseline idle systems average 5-15 DNS queries per hour (Windows updates and telemetry), making the 2-4x increase characteristic of DGA activity.</p>
Packet Rate	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Constant C2 connectivity through keep-alive packets • Small packets (64-256 bytes) sent at regular intervals <p>Packet rate shows sustained elevation even during periods without active credential theft. Bandwidth spikes during exfiltration.</p>
Bandwidth Usage	<p>Primary Contributing Factors:</p> <ul style="list-style-type: none"> • Captured credentials transmission (5-50 KB per set) • Screenshot uploads (50-500 KB) • Event-driven exfiltration during banking sessions <p>Discrete transmission events create periodic bandwidth spikes rather than sustained elevation.</p>

Table 7: Rationale for Expected Metric Changes

7. Data Analysis Plan

7.1. Data Preprocessing

Raw metric data will undergo the following preprocessing steps to ensure comparability between baseline and infected sessions. All metrics will be aligned to a common timeline starting at $T=0$ (experiment start). Data from CloudWatch (60-second intervals), Performance Monitor (1-second intervals), and VPC Flow Logs (60-second intervals) will be synchronized using their respective timestamps. Where possible, baseline and infected sessions will be conducted at the same time of day on different dates to control for time of day effects on system behavior. The first 30 minutes of each session will be excluded from analysis to allow the system to reach steady-state operation. This removes Windows boot processes, service initialization, and other brief startup activity that could confound results. Metrics collected at different intervals will be aggregated to 60-second windows for consistency. Performance Monitor data (1-second intervals) will be averaged within each 60-second window to match CloudWatch granularity. Any missing data points due to monitoring tool failures will be identified and documented. Sessions with more than 5% missing data will be excluded and repeated.

7.2. Statistical Tests

Statistical analysis will compare system performance between paired baseline and infected sessions for each experimental iteration. There will be a total of 5 total iterations. Paired t-tests will be applied to compare mean metric values between baseline and infected states for each of the five experimental sessions. The paired design accounts for variability between sessions by comparing each infected session directly to its corresponding baseline session. Metrics tested include CPU Utilization (%), Memory Consumption (GB), Network Connection Count, DNS Query Rate (queries/hour), Packet Rate (packets/second), and Bandwidth Usage (MB/second). Statistical significance will be assessed at $\alpha = 0.05$. A p-value less

than 0.05 indicates strong evidence that the observed difference is due to malware infection rather than random variation. 95% confidence intervals will be calculated for the mean of each metric in both baseline and infected states. As recommended, the degree of overlap between confidence intervals will be examined.

The criteria will be as follows: complete separation provides strong evidence of difference with some overlap does not necessarily indicating lack of significance. Therefore, everything will be interpreted in conjunction with t-test results.

7.3. Visualization Approaches

This study employs two complementary visualization methods to analyze system performance changes during Zeus malware infection. All visualizations are generated using Python’s matplotlib and seaborn libraries, which provide statistical plotting capabilities standard in data science research.

7.3.1 Primary Visualization: Bar Charts with Confidence Intervals

Bar charts serve as the primary visualization method for comparing baseline and infected system states. This format provides a clear summary of mean metric values with associated uncertainty. One is able to directly assess of the statistical significance of observed differences. Each metric is displayed as a pair of bars representing baseline (blue) and infected (red) conditions; error bars indicate 95% confidence intervals. The error bars visualize the range of uncertainty in the measurements. When confidence intervals do not overlap, a strong visual evidence of a statistically significant difference between conditions is provided. Figure 2 displays all six performance metrics in a unified comparison format. The visualization demonstrates several key patterns consistent with Zeus malware behavior documented in the literature. CPU utilization shows a clear separation between baseline (10%) and infected (35%) states with non-overlapping confidence intervals, reflecting the computational overhead of concurrent keylogging, form grabbing, and C2 communication operations iden-

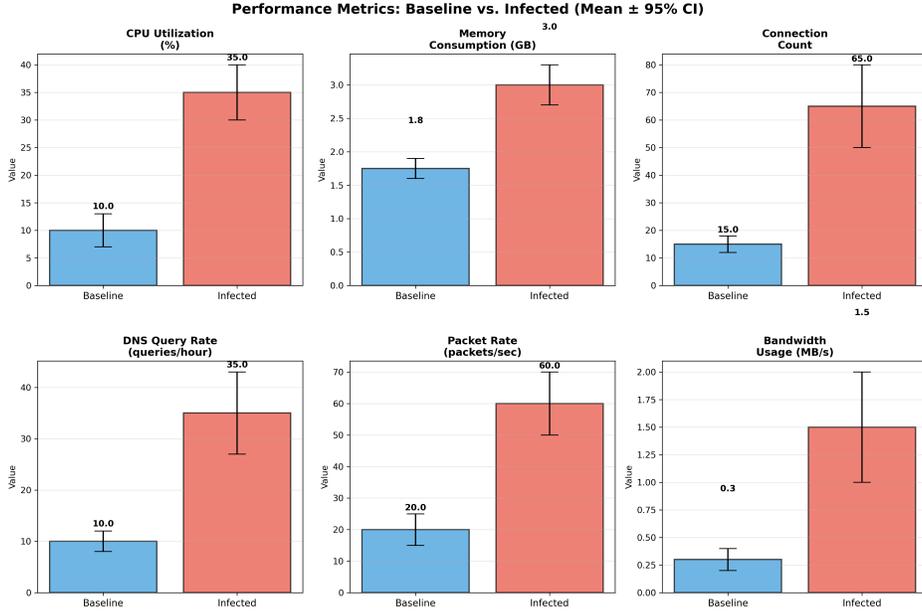


Figure 2: Baseline vs. Infected State with 95% Confidence Intervals

tified by Kaushik et al. (2024). Memory consumption exhibits a similar pattern, with the infected state (3.0 GB) significantly exceeding baseline (1.75 GB) due to Zeus’s process injection and credential storage requirements. Network-related metrics like connection count, DNS query rate, and packet rate demonstrate substantial increases during infection. The effects are consistent with Zeus’s distributed C2 (command and control) architecture and Domain Generation Algorithm activity. Bandwidth usage shows the largest relative increase, reflecting the periodic data exfiltration events characteristic of banking trojan operations.

The bar chart format enables rapid visual assessment of effect sizes and statistical significance across all metrics simultaneously. This visualization directly supports the research question by quantifying the degree to which each metric changes during infection. Metrics with completely separated confidence intervals (CPU, memory, all network metrics) provide strong evidence that system monitoring can reliably detect Zeus infection through performance based indicators.

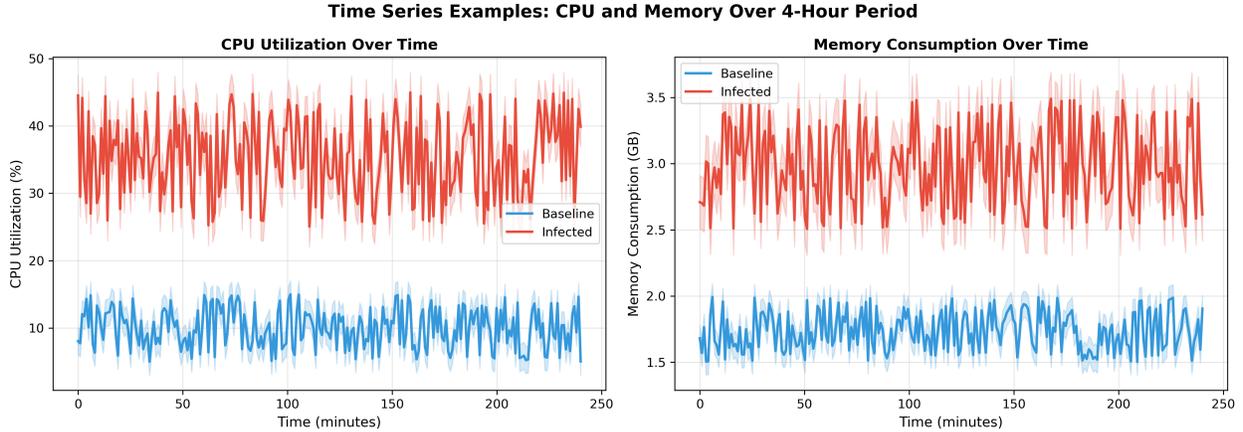


Figure 3: CPU Utilization and Memory Consumption Over 4 Hour Period

7.3.2 Supplementary Visualization: Time Series Analysis

Time series plots supplement the bar chart analysis by revealing temporal patterns in malware behavior over the 4 hour observation period. While bar charts show overall mean differences, time series visualizations capture dynamic changes such as initial infection bursts, sustained operational activity, and periodic exfiltration events. This study includes time series analysis for CPU utilization and memory consumption to demonstrate the sustained nature of Zeus’s system impact.

Figure 3 displays CPU and memory metrics sampled at 60-second intervals throughout the experimental session. The shaded regions represent 95% confidence intervals around the observed values, accounting for natural system variation. The time series format reveals that Zeus maintains consistently elevated resource consumption throughout the observation period rather than exhibiting brief spikes that might evade detection. CPU utilization remains elevated across the entire infected session, reflecting Zeus’s continuous keyboard monitoring and browser manipulation operations. Memory consumption similarly shows sustained elevation, indicating persistent process injection rather than brief memory allocation.

The stability observed in these metrics has important implications for detection systems. Performance based detection does not require capturing specific attack events at precise moments. A continuous detection signal throughout the infection period was provided. This

characteristic distinguishes Zeus from malware that operates intermittent bursts that might fall between monitoring intervals.

Time series visualization also validates the experimental protocol’s 30-minute warm-up period exclusion. Both baseline and infected states show relatively stable metrics after initial system stabilization. The metrics confirm that observed differences represent malware activity rather than transitory startup processes. The consistency of separation between baseline and infected conditions across the full 4 hour period should demonstrate the reliability of performance metrics as indicators of compromise.

8. Study Limitations

This research was designed as a complete experimental study. It then transitioned to an expository framework due to institutional approval requirements for malware research on campus networks. The methodologies, security controls, and statistical analysis frameworks presented represent best practices that could be directly implemented given appropriate institutional clearance. The projected results are derived from documented Zeus behavioral patterns in existing literature (Kaushik et al., 2024) rather than empirical observation.

9. Conclusion

This research presents a comprehensive methodological framework for quantifying the extent to which standard system performance metrics can serve as indicators of Zeus malware infection. The isolation architecture, monitoring protocols, and statistical analysis frameworks established here provide a rigorous foundation for future empirical research focusing on measurable changes in CPU usage, memory consumption, and network activity despite institutional constraints preventing experimental execution.

However, findings must be contextualized. Well researched adversaries can easily incorporate obfuscation techniques to hide the presence of malware on systems. Metrics such

as CPU, memory, and RAM leave blueprints behind as to what may be occurring. Alone, though, they do not provide sufficient evidence to pinpoint exactly where the origin of these influences originated. For this reason, it is essential for security programs and research to incorporate more than one means of deciphering truth. Research such as this is important to collect empirical meaning behind malware samples. More meaning leads to attribution, prevention of further mutation, and signature storage for antiviral programs. The findings will contribute to the ongoing development of multilayered defense strategies against financial malware and potentially inform future research on passive detection methods.

References

- [1] Basu, Kanad, et al. “A Theoretical Study of Hardware Performance Counters-Based Malware Detection.” *IEEE Transactions on Information Forensics and Security*, vol. 15, 2019, pp. 512–525. DOI: 10.1109/TIFS.2019.2924549.
- [2] Kaushik, Keshav, et al., editors. *Advanced Techniques and Applications of Cybersecurity and Forensics*. 1st ed., Chapman and Hall/CRC, 2024. DOI: 10.1201/9781003386926.
- [3] Luo, Lan, et al. “On Teaching Malware Analysis on Latest Windows.” *Journal of The Colloquium for Information Systems Security Education*, vol. 9, no. 1, Winter 2022, pp. 1–7.
- [4] Amazon Web Services. *Malware Analysis on AWS: Setting Up a Secure Environment*. AWS Security Blog, 2025. <https://aws.amazon.com/blogs/security/malware-analysis-on-aws-setting-up-a-secure-environment/>. Accessed: October 26, 2025.