A Purple Team Security Assessment of a Maturity Level 1 Healthcare Provider: Detection Engineering and Validation on a Segmented Enterprise Network

Final Project Report

Group 2

, Abdul-Ramon Adeyemi Folarin, &

, , , &

School of Information Technology Administration & Security, Seneca Polytechnic

SPR708NBB: Attack and Defense

Professor Professor Ahmad Ghobeiti

August 12, 2026

Student Integrity Statement: I/we declare that the attached assignment is my/our own work in accordance with the Seneca Academic Policy. No part of this assignment has been copied manually or electronically from any other source (including web sites) or distributed to other students

Executive Summary

This report documents a purple team security assessment conducted against a segmented enterprise network modelling a mid-size Ontario healthcare provider assessed at Security Maturity Level 1. The exercise follows the Atomic Purple Team lifecycle and the SCYTHE Purple Team Exercise Framework, with adversary behaviour mapped to MITRE ATT&CK version 19 and risk framed using NIST SP 800-30. A single connected intrusion was executed from an external attacker, through an exposed web application, to a domain-wide compromise ending in a simulated ransomware event, while a defensive team recorded prevention, alerting, and logging outcomes at each step. The baseline pass exposed critical weaknesses, including chained web-application vulnerabilities enabling remote code execution, credentials stored in files, a domain administrator compromise through credential reuse, and an absence of endpoint telemetry on the most sensitive hosts. Four control gaps were remediated: a web application firewall was placed in front of the vulnerable application, file integrity monitoring was extended to the file server, the SIEM host was restricted and given a self-protection rule, and a backup and recovery capability was established. A validation pass confirmed measurable improvement, with the initial exploit blocked, file changes detected, lateral movement to the SIEM alerted, and the ransomware impact reversed from backup. The assessment demonstrates a repeatable and honestly documented path from Maturity Level 1 toward improved detection and recovery.

Table of Contents

Executive Summary 2
Table of Figures 4
Table of Tables 5
1 Introduction 6
2 Market Segment Analysis 6
2.1 Current Threat Landscape 6
2.2 Threat Actors and Techniques 6
2.3 Relevant Threat Intelligence and Risk Assessment 7
2.4 Regulatory Environment 7
2.5 Influence on Exercise Design 7
3 Methodology and Frameworks 7
4 Target Environment 8
5 Exercise Execution (Baseline) 9
5.1 Preparation and Baseline Monitoring 10
5.2 Reconnaissance 12
5.3 Initial Access, Execution, and Persistence 14
5.4 Credential Access 16
5.5 Lateral Movement to the SIEM 17
5.6 Defense Impairment 18
5.7 Domain Compromise 18
5.8 Impact 20
6 Detection Analysis and Findings 20
7 Remediation 22
7.1 File Server Telemetry 22
7.2 Web Application Firewall 23
7.3 SIEM Hardening 23
7.4 Backup and Recovery 24
7.5 Domain Controller Telemetry 24
8 Validation 25
9 Discussion and Lessons Learned 27
10 Conclusion 27
References 28

Table of Figures

Figure 1 Target enterprise topology 8
Figure 2 Central monitoring confirmed on the SIEM host 10
Figure 3 Registered agent inventory 10
Figure 4 Wazuh dashboard agent overview panel 11
Figure 5 Network intrusion detection sensor active on the core router 11
Figure 6 Domain Controller agent active before the exercise 12
Figure 7 Domain Controller agent status confirmation 12
Figure 8 Reconnaissance scan against the file server 13
Figure 9 Network IDS produced no alert for the scan 13
Figure 10 Nikto web application scan output against the file server 14
Figure 11 Web application exploited to achieve code execution 14
Figure 12 Web shell file confirmed on the host 15
Figure 13 No prevention or detection of the foothold in the baseline 15
Figure 14 Interactive command execution and persistence 16
Figure 15 Cron persistence confirmed on the file server 16
Figure 16 Credentials recovered from the host 17
Figure 17 Lateral movement to the SIEM challenged by multi-factor authentication 17
Figure 18 SIEM does not monitor its own authentication log 18
Figure 19 Defensive monitoring disabled from the SIEM host 18
Figure 20 Domain Controller compromised and enumerated 19
Figure 21 Domain Controller agent collecting no Windows telemetry 19
Figure 22 Empty Wazuh event query for the Domain Controller agent 20
Figure 23 Ransomware simulation against the file share 20
Figure 24 File integrity monitoring confirmed on the file server 22
Figure 25 File server agent active with monitoring applied 23
Figure 26 Web application firewall blocking the exploit 23
Figure 27 SIEM self-protection rule validated 24
Figure 28 Baseline backup established 24
Figure 29 Exploit blocked on the validation re-run 25
Figure 30 ModSecurity WAF setup and rule configuration 25
Figure 31 File change now detected by the SIEM 26
Figure 32 Lateral movement to the SIEM blocked and alertable 26
Figure 33 Ransomware impact reversed from backup 26

Table of Tables

Table 1 Systems in scope and their role in the exercise 8
Table 2 Attack chain mapped to MITRE ATT&CK version 19 9
Table 3 Baseline detection outcomes per technique (P prevented, A alerted, L logged) 20
Table 4 Findings and recommendations. 21
Table 5 Before-and-after validation summary 25

1 Introduction

Purple teaming is a collaborative security exercise in which an offensive team emulates realistic adversary behaviour while a defensive team measures, in real time, whether each action is prevented, detected, or merely logged. Rather than treating attack and defence as separate engagements, the purple model runs them together so that every technique produces an immediate, testable answer about the organisation's visibility. This report applies that model to a segmented enterprise network that represents the internal infrastructure of a mid-size healthcare provider operating at Security Maturity Level 1.

The objective of the exercise was to determine whether such an organisation can detect and respond to a realistic intrusion that culminates in ransomware, and to measure the improvement achieved after targeted remediation. A single connected attack chain was executed from an external attacker to full domain compromise, the detection outcome of every step was recorded, four control gaps were remediated, and the same chain was re-run to validate the improvement. Throughout, findings are reported honestly: steps that were blocked, telemetry that was missing, and controls that behaved unexpectedly are documented as they occurred, because an accurate account of gaps is the primary value a purple team delivers.

The remainder of the report is organised as follows. Section 2 analyses the healthcare market segment and its threat landscape. Section 3 describes the methodology and frameworks used. Section 4 describes the target environment and its maturity characteristics. Section 5 presents the baseline execution of the attack chain. Section 6 analyses detection outcomes and enumerates findings. Section 7 details the remediation of each gap, and Section 8 presents the validation pass. Section 9 discusses lessons learned, and Section 10 concludes.

2 Market Segment Analysis

Healthcare was selected as the market segment because it combines a high-value target profile with the resource constraints that characterise Maturity Level 1, and because its regulatory environment makes the consequences of the modelled attack concrete. The justification below addresses the current threat landscape, the active threat actors and their techniques, relevant threat intelligence, the regulatory environment, and the resulting risk assessment, and then explains how each shaped the exercise design.

2.1 Current Threat Landscape

Healthcare is consistently the most targeted critical-infrastructure sector. Industry reporting indicates that ransomware attacks on healthcare rose by roughly fourteen percent in the first half of 2026 to about 410 incidents worldwide, an average of more than two per day, and that Canada ranked among the five most affected countries alongside the United States, Germany, India, and Australia. United States federal reporting similarly identified healthcare as the most targeted critical-infrastructure subsector in its most recent annual crime figures. Attacks increasingly reach providers not only directly but through the vendors and third parties in the healthcare supply chain, widening the effective attack surface of any single organisation.

2.2 Threat Actors and Techniques

The groups most active against healthcare providers during this period were Qilin, The Gentlemen, LockBit, and INC, with Qilin the single most active against providers. These are predominantly financially motivated, frequently Russian-speaking, double-extortion operations that exfiltrate data before encrypting it. Their tradecraft follows a consistent pattern that this exercise deliberately mirrors: initial access through an exposed application or phishing, credential theft, lateral movement across a flat internal network, and finally data encryption for impact. The median time from first access to encryption in healthcare is now reported at four to five days, while the average time to identify and contain a healthcare breach is around

279 days, which is precisely why detection speed, and not only prevention, is the metric that matters for this sector.

2.3 Relevant Threat Intelligence and Risk Assessment

Average ransom demands in confirmed healthcare incidents rose sharply, with first-quarter figures for 2026 reported near 16.9 million dollars. Canadian incidents anchor the risk locally: a 2023 ransomware attack on a shared-services provider disrupted several southwestern Ontario hospitals through a trusted third party, and a subsequent privacy ruling reinforced the regulatory stakes. Applying the risk model of NIST SP 800-30, the combination of high-value personal health information, high operational urgency that raises willingness to pay, and Maturity Level 1 controls yields a high likelihood and a high impact rating. This assessment justifies prioritising the four control gaps that this exercise closes.

2.4 Regulatory Environment

Ontario healthcare organisations are bound by the Personal Health Information Protection Act, which requires notification to affected individuals at the first reasonable opportunity, interpreted by the provincial Information and Privacy Commissioner as approximately seventy-two hours, far tighter than the sixty-day window of the United States HIPAA framework. Following a 2025 ruling upheld on appeal, ransomware that only encrypts data is a reportable breach even without confirmed exfiltration, and recent amendments introduced mandatory regulator notification and the first administrative monetary penalties, reported up to 500,000 dollars for organisations. The federal Personal Information Protection and Electronic Documents Act provides a further backstop.

2.5 Influence on Exercise Design

Because a mid-size clinic cannot absorb multi-day attacker dwell time, the exercise measures detection and recovery speed rather than prevention alone. Because encryption by itself is a reportable breach, the impact stage was included even as a simulation. Because the sector is compromised through exposed applications and flat networks, the attack chain begins at the web and file server and moves laterally, and the remediation targets endpoint telemetry, web application firewalling, SIEM isolation, and backup.

3 Methodology and Frameworks

The exercise was conducted using the Atomic Purple Team lifecycle, an iterative cycle of plan, attack, detect, and defend, in combination with the SCYTHE Purple Team Exercise Framework. The framework prescribes that each technique be emulated one at a time while the offensive and defensive teams jointly answer three questions: was the activity prevented, was it alerted, and was it logged at all. Where a technique passed undetected, the defensive team tuned a detection and the offensive team repeated the technique to confirm the fix. This per-technique loop is the mechanism by which detection engineering was performed and validated.

Adversary behaviour was mapped to MITRE ATT&CK version 19. A notable change in this version is that the former Defense Evasion tactic has been divided into two tactics, Stealth and Defense Impairment; the step in which the attacker disabled monitoring is therefore mapped to Defense Impairment. Risk was framed using NIST SP 800-30, and the incident-handling perspective drew on NIST SP 800-61. Roles were assigned across the four team members covering offensive execution, defensive monitoring and detection engineering, purple coordination and ATT&CK mapping, and threat intelligence with reporting.

4 Target Environment

The exercise ran inside a segmented enterprise network that models the internal infrastructure of a mid-size healthcare provider. The environment runs an Active Directory domain, a file server used to exchange patient-related documents, a network intrusion detection sensor, and a centralised SIEM, but it lacks a dedicated security operations function and mature incident response, which places it at Maturity Level 1. Maturity is a property of operations and process, not of scale: the organisation owns capable tooling, yet that tooling is incompletely configured, and it is these gaps, rather than any absence of technology, that the exercise exposes.

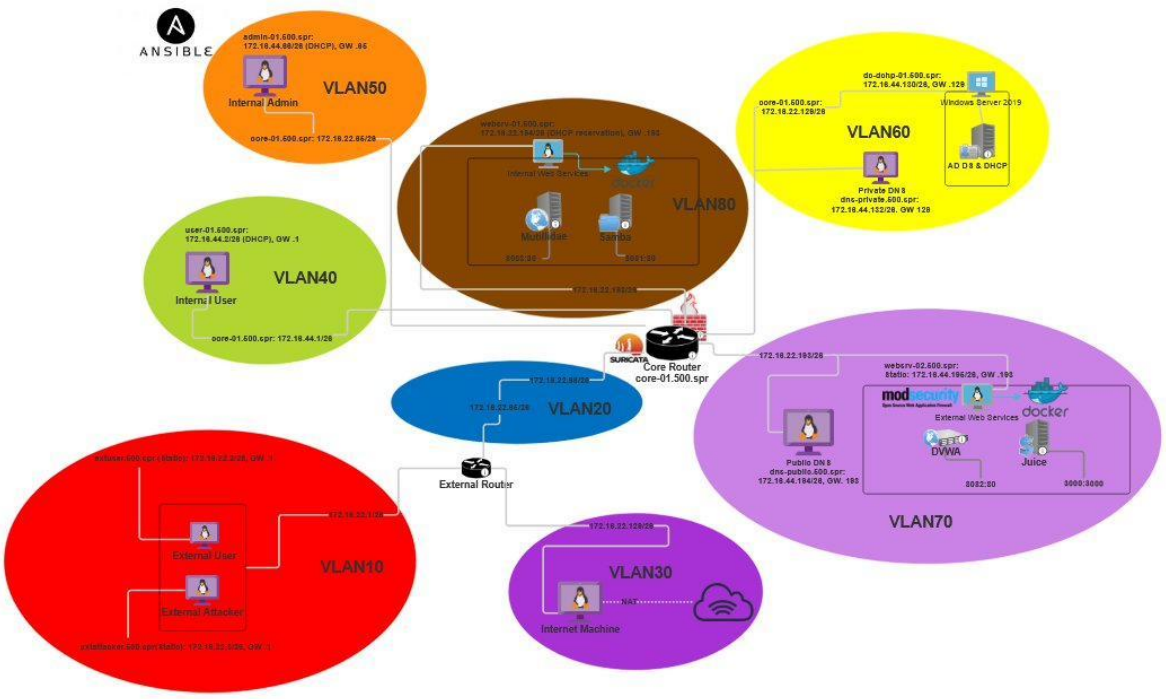


Figure 1 Target enterprise topology

Figure 1 shows a segmented multi-VLAN network that models the healthcare provider, separating external, web and file, user, management, and network-services zones. Five systems were placed in scope for the exercise: the external attacker, the file server, the core router and IDS sensor, the SIEM, and the Domain Controller. The layout provides the multi-node complexity required to exercise a realistic intrusion while retaining clear segmentation boundaries. The remaining hosts were held out of scope to keep the assessment focused on the modelled attack path.

Table 1 Systems in scope and their role in the exercise

System	Role in the exercise	Zone
External attacker	Offensive host executing the full attack chain	External VLAN
File server	Exposed web application and file share; first target	Private apps VLAN
Core router / IDS	Routing and network intrusion detection sensor	Core / transit

System	Role in the exercise	Zone
SIEM	Centralised monitoring; also an SSH lateral target	Management VLAN
Domain Controller	Active Directory and DHCP; domain lateral target	Network services VLAN

Four documented weaknesses defined the Maturity Level 1 baseline and became the focus of remediation: the file server had no file integrity monitoring covering the attack surface; the vulnerable web application was not protected by a web application firewall; the SIEM host was reachable over SSH from an untrusted segment; and no pre-incident backup or recovery process existed. Each was left open for the baseline pass and closed during remediation.

5 Exercise Execution (Baseline)

The attack chain was executed as a single connected intrusion so that detection gaps would compound realistically and map directly to the environment's documented weaknesses. Table 2 summarises the techniques and their mapping to MITRE ATT&CK version 19. Each subsection below narrates one stage and presents the corresponding evidence.

Table 2 Attack chain mapped to MITRE ATT&CK version 19

Step	Tactic (v19)	Technique	Outcome (baseline)
1	Reconnaissance	T1595 Active Scanning	Not detected by IDS
2	Initial Access	T1190 Exploit Public-Facing Application	Succeeded, not blocked
3	Execution / Persistence	T1059; T1505.003 Web Shell; T1053.003 Cron	Succeeded, no alert
4	Credential Access	T1552.001 Credentials in Files	Credentials recovered
5	Lateral Movement	T1021.004 SSH; T1078 Valid Accounts	Reached host, MFA blocked
6	Defense Impairment	T1562.001 Impair Defenses	Monitoring disabled
7	Lateral Movement	T1021.002 SMB; T1110.003 Password Spraying	Domain admin, no telemetry
8	Impact	T1486 Data Encrypted for Impact (simulation)	No alert, no recovery

5.1 Preparation and Baseline Monitoring

Before the chain was executed, the monitoring stack was confirmed operational so that later detection gaps could be attributed to configuration rather than to absent infrastructure. The SIEM and its supporting components were verified running, and the agent inventory was recorded.

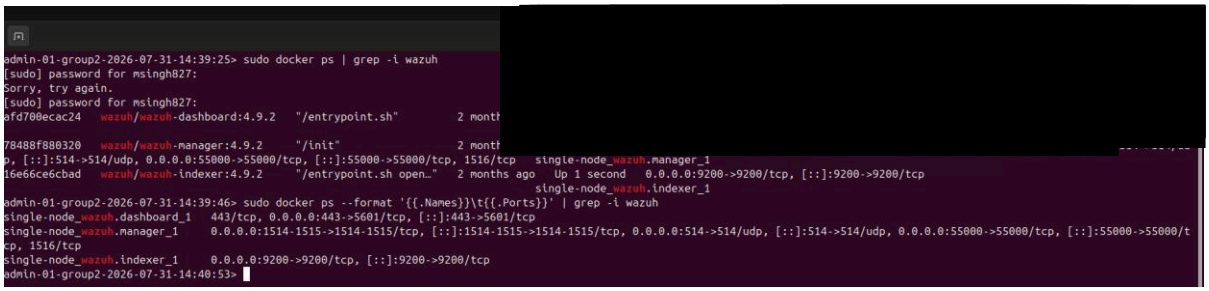


Figure 2 Central monitoring confirmed on the SIEM host

Figure 2 confirms that the SIEM manager, indexer, and dashboard components are running as a single-node containerised deployment before the exercise begins. It identifies the published ports used for agent enrolment and analyst access, establishing that centralised monitoring is available. Confirming this baseline is a prerequisite for measuring which attacker actions are subsequently detected. It is consistent with a Maturity Level 1 organisation that owns tooling but has not fully operationalised it.

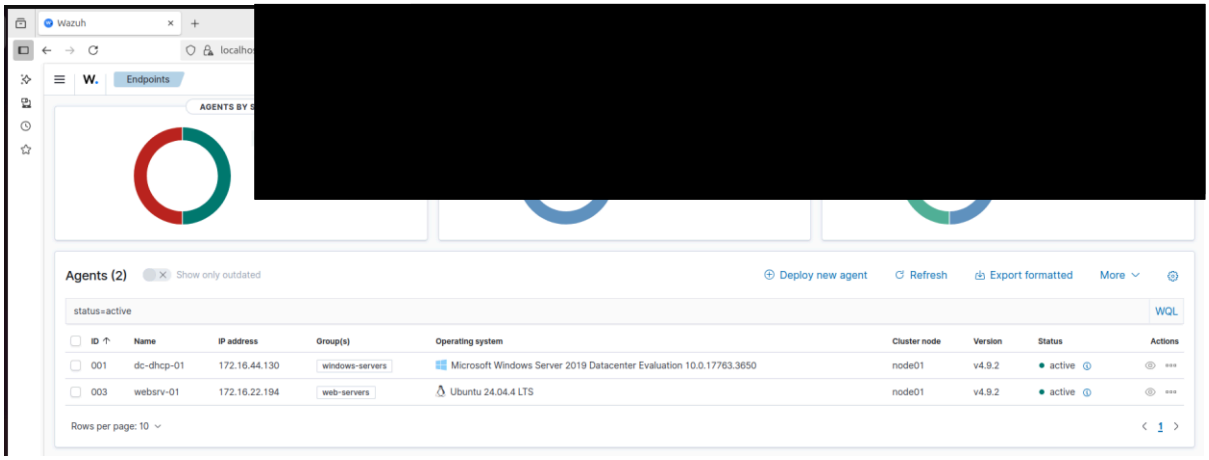


Figure 3 Registered agent inventory

Figure 3 shows the listing shows the endpoints enrolled with the SIEM, with the Domain Controller and file server active and two secondary hosts disconnected. This inventory records which systems were reporting telemetry at the start of the exercise. Establishing the agent state up front allows later detection gaps to be attributed to configuration rather than to a missing agent. As the exercise shows, active agents that nonetheless fail to detect key activity become a central finding.

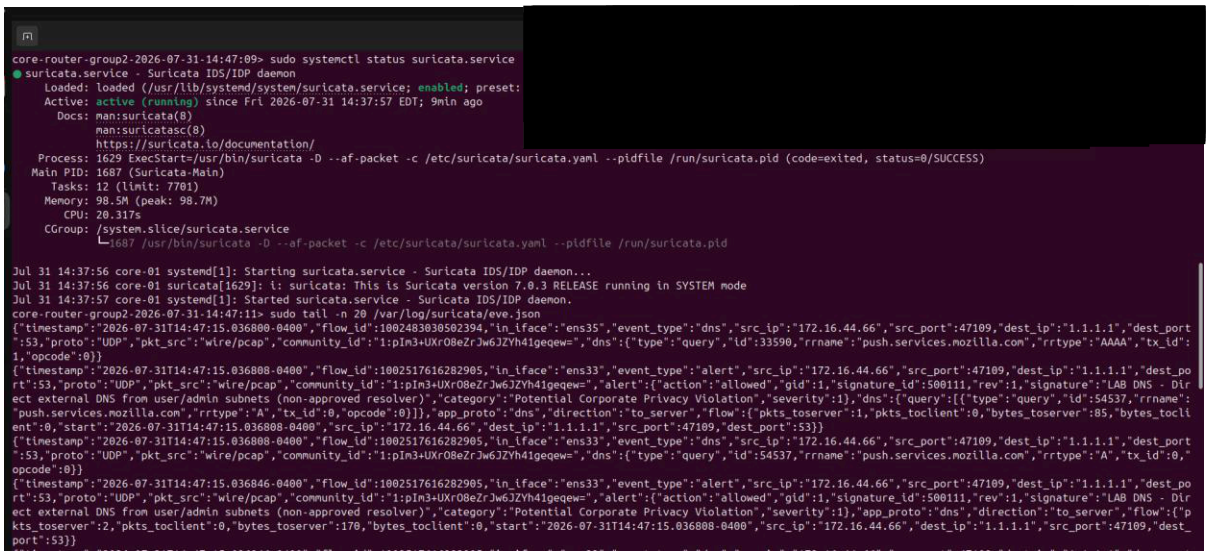


Figure 4 Wazuh dashboard agent overview panel

Figure 4 shows the dashboard Agents view confirms the enrolled endpoints and their status at the start of the exercise, providing a visual counterpart to the command-line agent listing. It shows which systems were actively reporting telemetry and highlights the two disconnected secondary hosts that were intentionally out of scope. This panel establishes the monitoring baseline before any attacker activity begins. It corroborates the agent-control output from the previous figure.

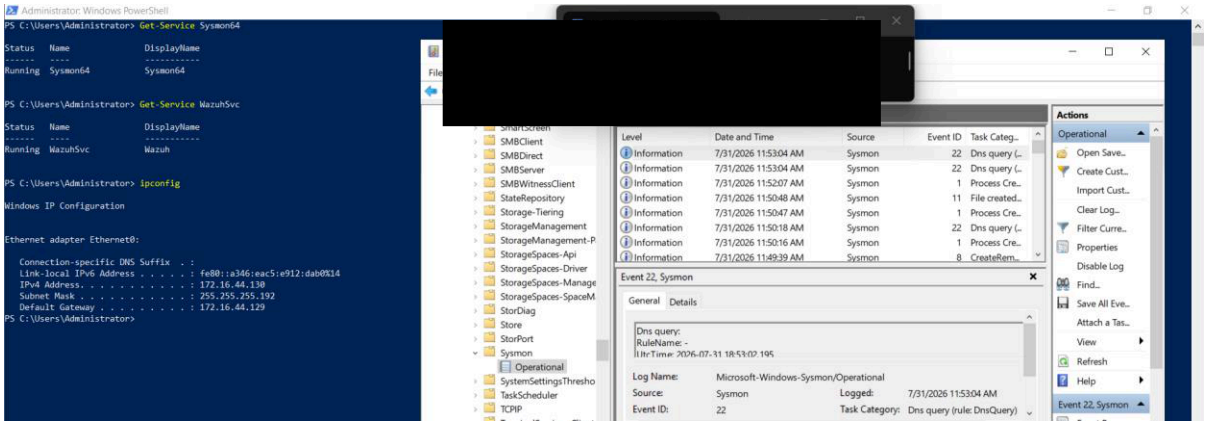


Figure 5 Network intrusion detection sensor active on the core router

Figure 5 shows the output confirms that the Suricata sensor is running on the core router and writing events before the exercise begins. It verifies that traffic between the attacker path and the internal zones is under inspection. Establishing that the network sensor is operational allows the later absence of a scan alert to be attributed to rule coverage rather than to an offline sensor. This forms part of the baseline instrumentation against which detection outcomes are measured.

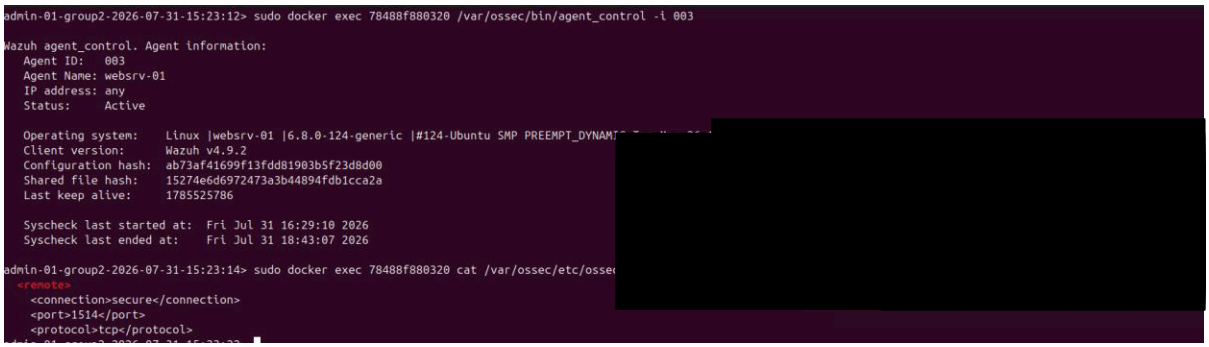


Figure 6 Domain Controller agent active before the exercise

Figure 6 shows the agent status confirms that the Domain Controller is enrolled with the SIEM and reporting at the start of the exercise. This establishes that any later absence of Windows telemetry is a configuration gap rather than a missing agent. It provides the baseline for the domain-side monitoring evaluated during the lateral-movement stage. The distinction becomes central to the finding that the enrolled agent collected no event data.

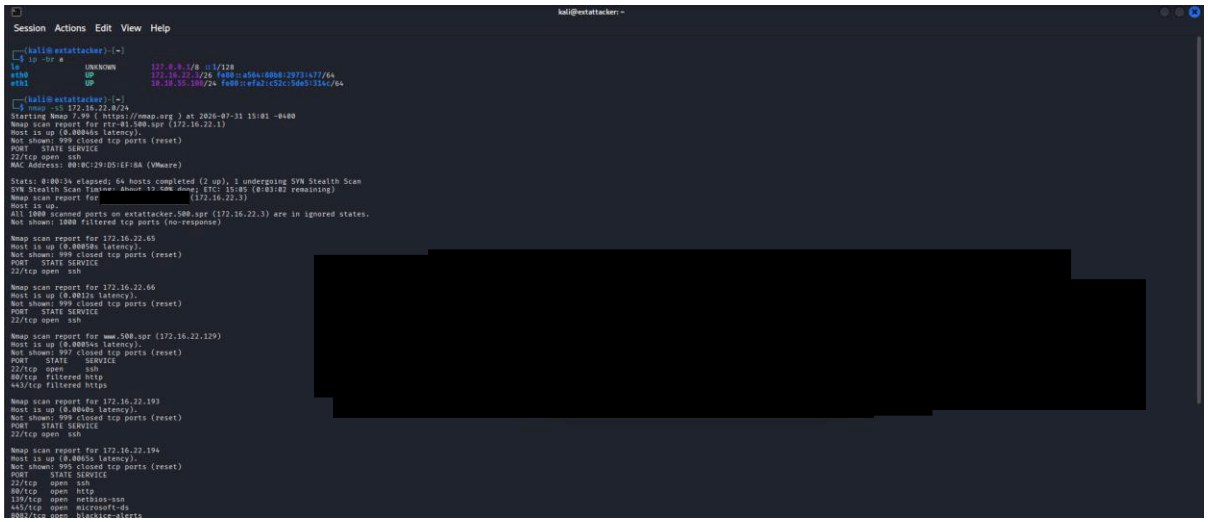


Figure 7 Domain Controller agent status confirmation

Figure 7 shows a second view of the Domain Controller agent confirms active status and the recent syscheck scan window, corroborating the preceding agent-information output. It makes clear the agent is functional and communicating rather than failing silently. The syscheck timestamps confirm the agent ran its file integrity pass during the exercise window. This paired evidence strengthens the finding that the absence of alerts was a configuration gap, not an infrastructure failure.

5.2 Reconnaissance

The attacker began by enumerating the file server from the external segment to identify the initial attack surface.



Figure 8 Reconnaissance scan against the file server

Figure 8 shows the scan reveals an exposed web application on a non-standard port, file-sharing services, and a second web service leaking a source-code repository and directory listings. This reconnaissance identifies the attack surface and directly informs the exploitation path selected next. The breadth of exposed services on a single host reflects the weak perimeter posture typical of the modelled maturity level. It provided several viable entry points before any exploitation was attempted.

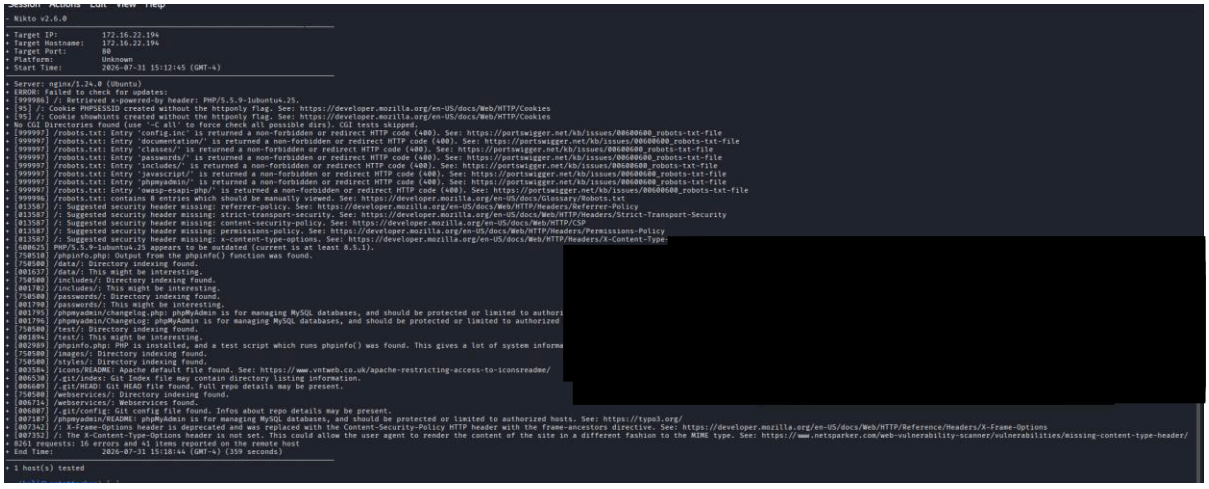


Figure 9 Network IDS produced no alert for the scan

Figure 9 shows **me** querying the network intrusion detection log for the attacker source address during the scanning window returns no scan-related alerts. This confirms the sensor was not loaded with a scan-detection rule, leaving an aggressive full-port sweep and an automated web crawl invisible to the monitoring stack. The absence is recorded as a detection gap rather than treated as a passing result. It illustrates that owning an IDS does not, by itself, provide coverage without tuned rules.

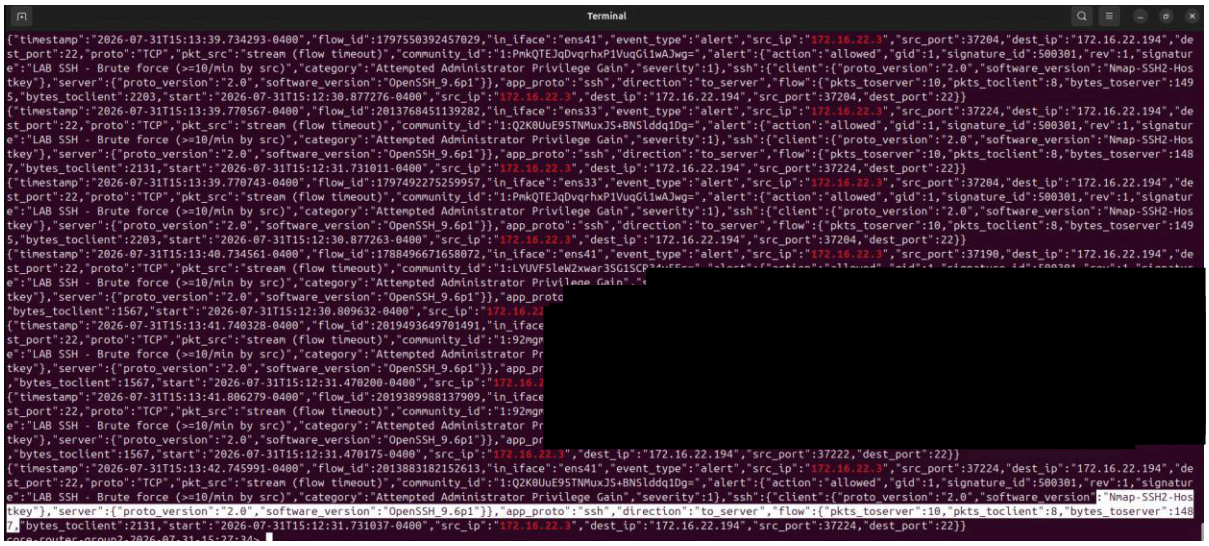


Figure 10 Nikto web application scan output against the file server

Figure 10 shows the Nikto output enumerates misconfigurations found during the web crawl, including missing security headers, an exposed phpinfo page, directory indexing, and a visible source-code repository. Each finding represents an additional potential entry point beyond the command-injection vector exploited in the next step. The depth of the Nikto output also illustrates why the network IDS should have raised an alert on the automated scan traffic. It supports the case for both a scan-detection rule and a WAF as complementary controls.

5.3 Initial Access, Execution, and Persistence

Exploitation of the web application produced remote code execution through a chain of two distinct weaknesses. Because the application web root was not writable by the web-server account, a minimal payload was staged in a temporary directory using a command-injection flaw and then executed indirectly through a separate file-inclusion vulnerability. This chained path is more representative of real exploitation than a single-step web-shell drop.

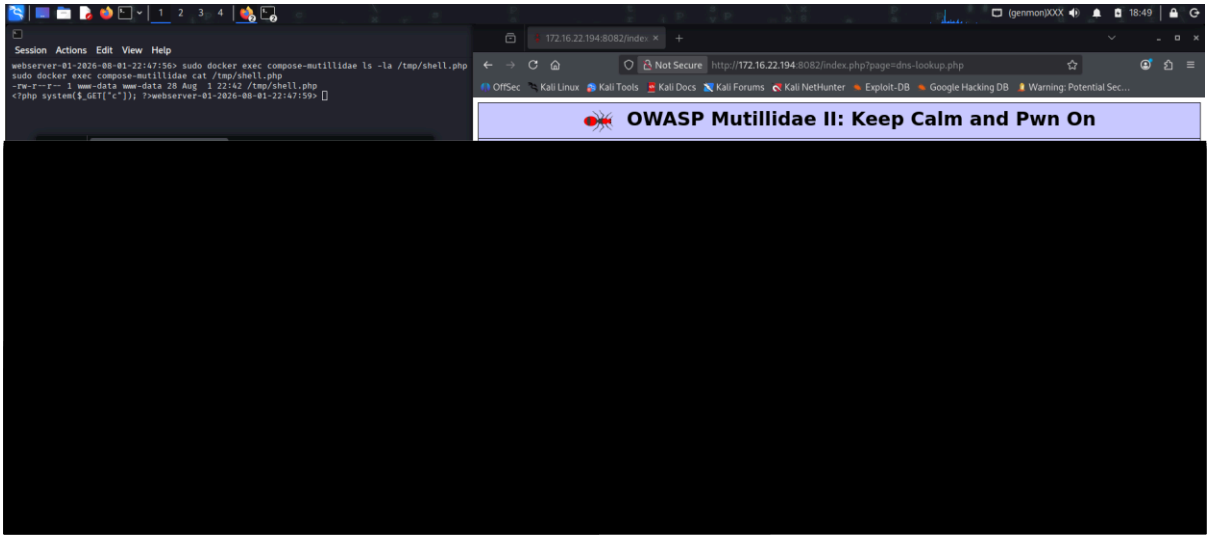


Figure 11 Web application exploited to achieve code execution

Figure 11 shows the evidence demonstrates a command-injection flaw used to write a minimal web shell, which is then executed through a separate file-inclusion vulnerability because the web root was not writable.

The returned identity confirms that commands run in the context of the web-server account. Chaining the two weaknesses produced reliable remote code execution on the file server. This established the initial foothold and maps to exploitation of a public-facing application.

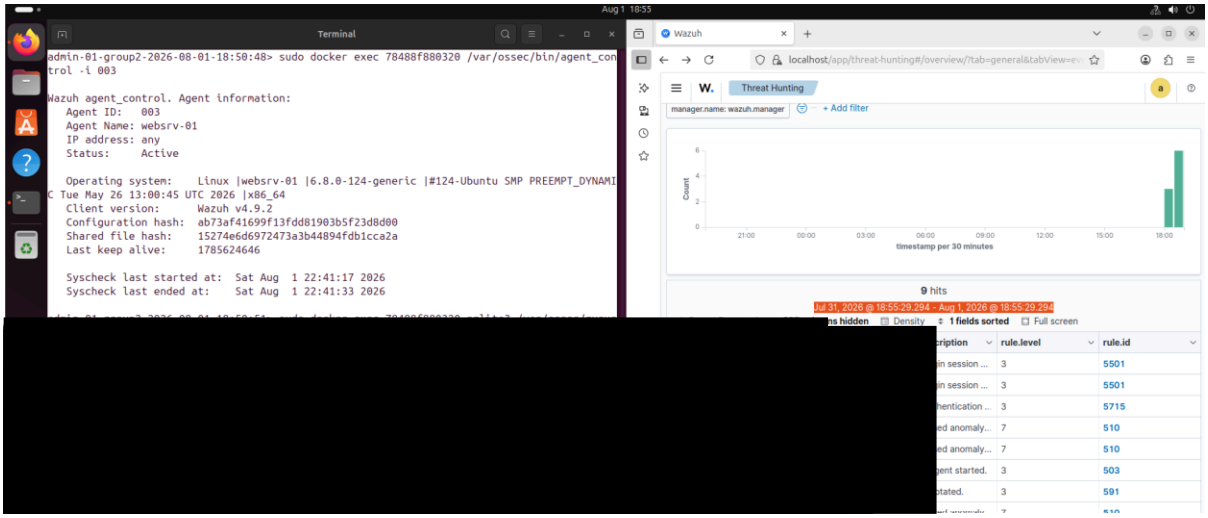


Figure 12 Web shell file confirmed on the host

Figure 12 shows the listing of the temporary directory and the cat of the shell file confirm that the PHP payload was successfully staged on the file server. It shows the file owner, permissions, and content, confirming the shell is ready to execute arbitrary commands. This supplements the browser-side execution evidence by showing the artefact from the server perspective. Together the two captures prove the full chain from injection to durable foothold.

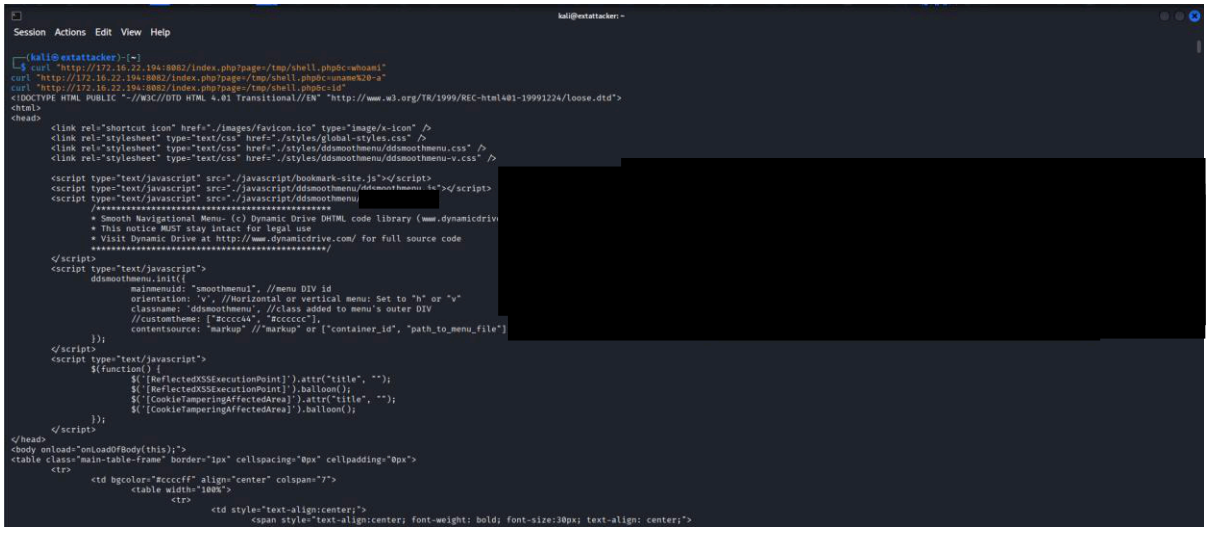


Figure 13 No prevention or detection of the foothold in the baseline

Figure 13 shows during the baseline pass the exploitation was neither blocked nor alerted on. The vulnerable application sat directly exposed with no web application firewall in its request path, and file integrity monitoring did not cover the location used to stage the payload. A successful remote code execution event therefore generated no signal in the SIEM. This combined gap in prevention and detection is closed and re-tested in later phases.

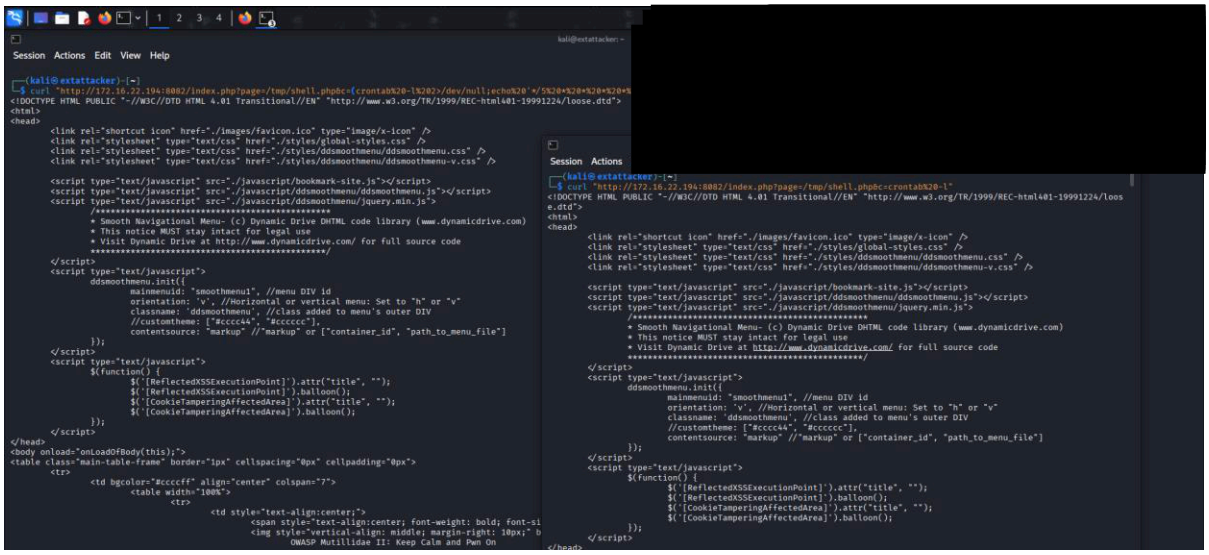


Figure 14 Interactive command execution and persistence

Figure 14 shows further commands issued through the web shell confirm interactive control of the host and establish persistence through a scheduled task. The output verifies the operating context and shows that the foothold can survive beyond the initial session. Persistence of this kind is behaviour that a tuned endpoint agent would be expected to flag. In the baseline environment the activity proceeded without any corresponding alert.

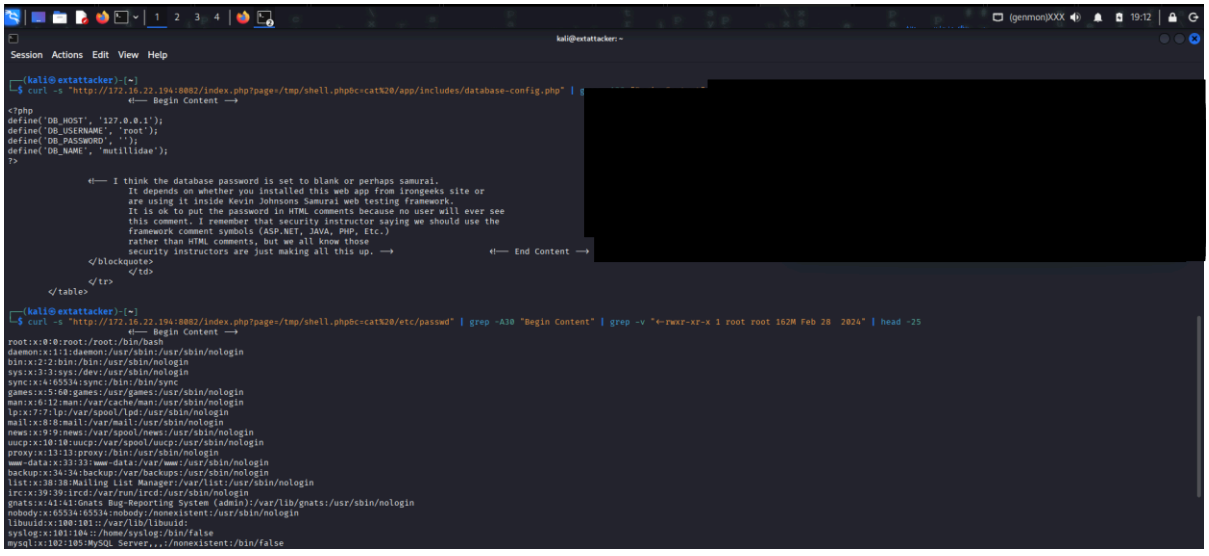


Figure 15 Cron persistence confirmed on the file server

Figure 15 shows a second terminal panel confirms the scheduled task entry was written to the crontab, verifying that the persistence mechanism survived beyond the initial shell session. It shows the attacker maintaining access independently of the original web-shell connection. Persistence of this kind would ordinarily trigger a rootcheck or process-creation alert in a mature environment. Its absence in the SIEM confirms that no scheduled-task monitoring was active during the baseline pass.

5.4 Credential Access

The foothold was used to search the host for reusable secrets that would enable lateral movement.

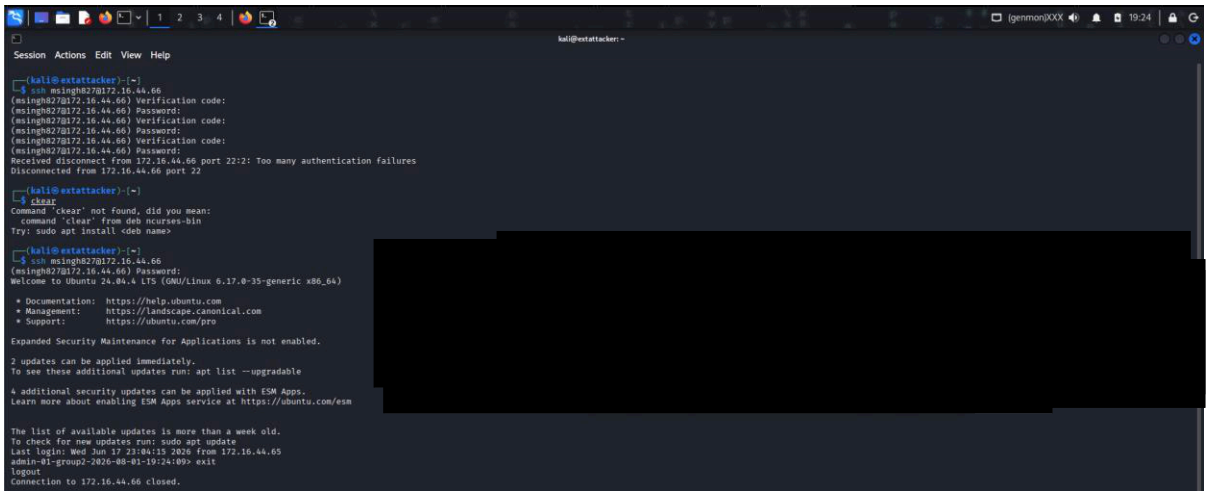


Figure 16 Credentials recovered from the host

Figure 16 shows a configuration file recovered through the web shell exposes a database account with a blank password, and an adjacent comment discloses a fallback credential in clear text. Storing secrets in application files and in client-visible comments is a direct instance of unsecured credentials. The recovered material, together with credentials harvested from the application, supplied the account later reused against the domain. This finding links the web compromise directly to the subsequent domain compromise.

5.5 Lateral Movement to the SIEM

The attacker attempted to move laterally onto the monitoring host itself, the strongest narrative point in the chain because the SIEM is both the security control point and a reachable target.

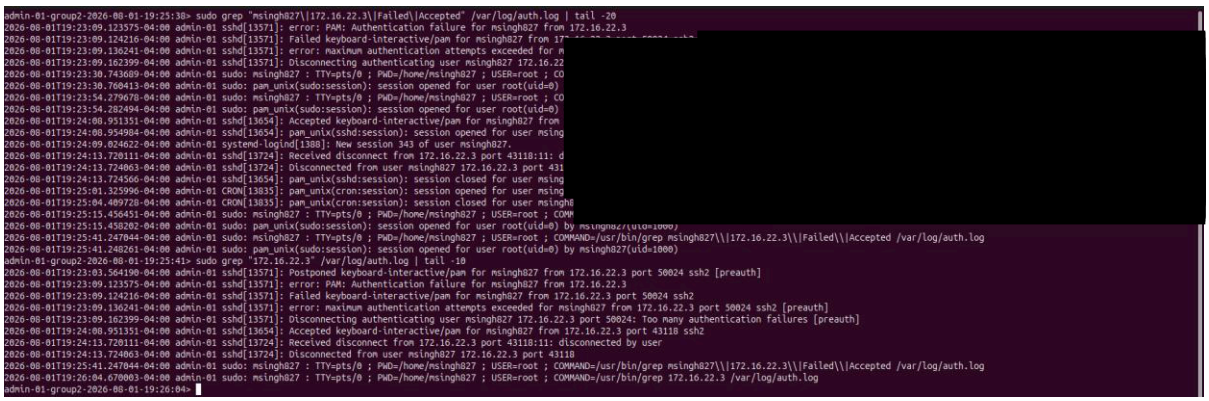


Figure 17 Lateral movement to the SIEM challenged by multi-factor authentication

Figure 17 shows an SSH attempt from the external attacker host to the SIEM reached the service and returned a multi-factor verification prompt before rejecting the login. This demonstrates two things at once: the monitoring host is reachable over SSH from an untrusted segment, and multi-factor authentication provides a genuine barrier to credential misuse. The network exposure remains a finding because the attack surface exists irrespective of whether authentication ultimately succeeds. It is a nuanced result showing a partial control rather than a clean success or failure.

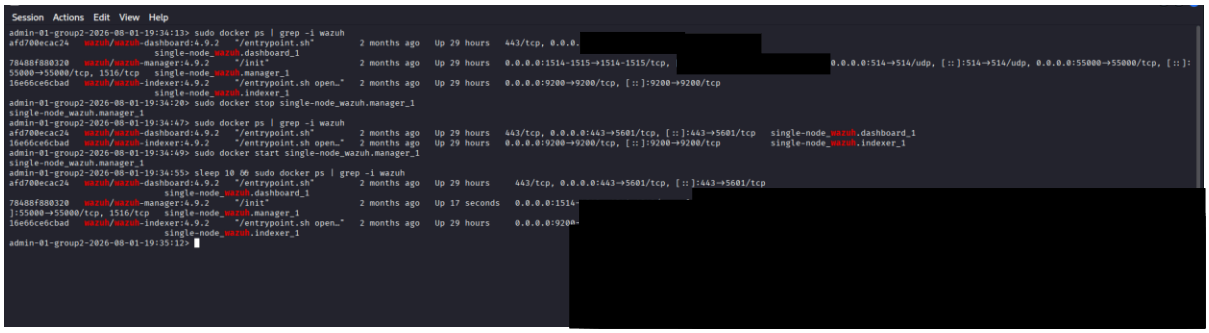


Figure 18 SIEM does not monitor its own authentication log

Figure 18 shows Authentication events on the SIEM host are recorded locally but are not ingested by the manager, which runs in a container without visibility of the host log. Login attempts against the monitoring platform itself therefore generate no alert. For a host that is simultaneously the monitoring point and a reachable target, this is a significant blind spot. The gap directly motivates the self-protection rule introduced during remediation.

5.6 Defense Impairment

From the position of an attacker who has reached the SIEM host, the central monitoring service was stopped and then restored, emulating impairment of defences.

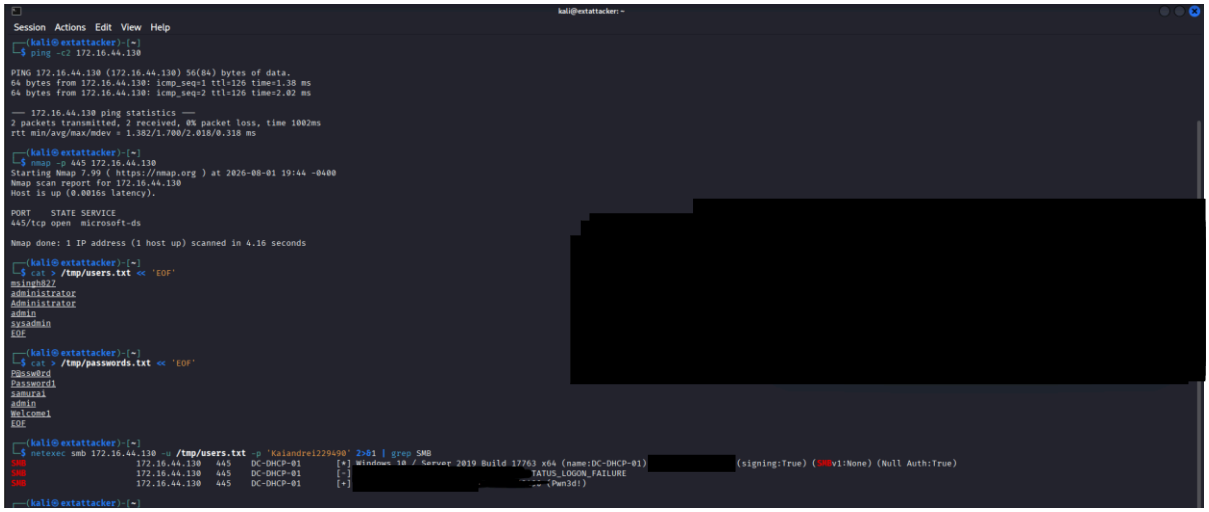


Figure 19 Defensive monitoring disabled from the SIEM host

Figure 19 shows the central monitoring service is stopped from the SIEM host and then restored, emulating an adversary impairing defences after gaining access, mapping to the Defense Impairment tactic in the current ATT&CK version. The step produced an unintended operational consequence: stopping the manager disconnected every enrolled agent, and the Windows Domain Controller agent did not automatically reconnect. This cascading loss of visibility is itself an important finding about the resilience of the monitoring architecture. It shows that impairing a single central component degraded coverage across the whole estate.

5.7 Domain Compromise

Credentials harvested from the compromised application were reused against the Domain Controller in a controlled password-spraying attempt.

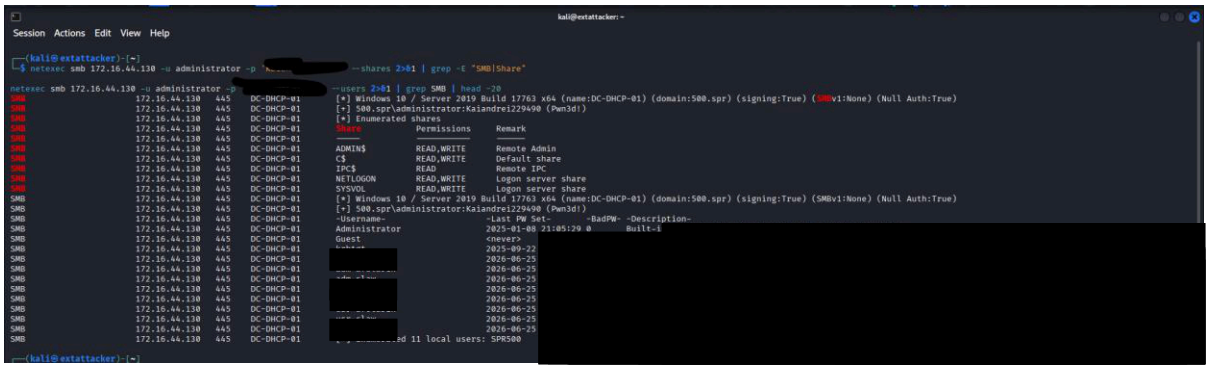


Figure 20 Domain Controller compromised and enumerated

Figure 20 shows a password-spraying attempt against the Domain Controller using a harvested credential results in a successful authentication as the domain administrator, with the tool confirming full administrative control. It enumerates the default administrative shares with read and write access and lists every domain account, including the administrative and standard accounts belonging to the team. This single credential-reuse finding demonstrates the blast radius achievable against a flat domain. It represents the most severe outcome of the exercise.

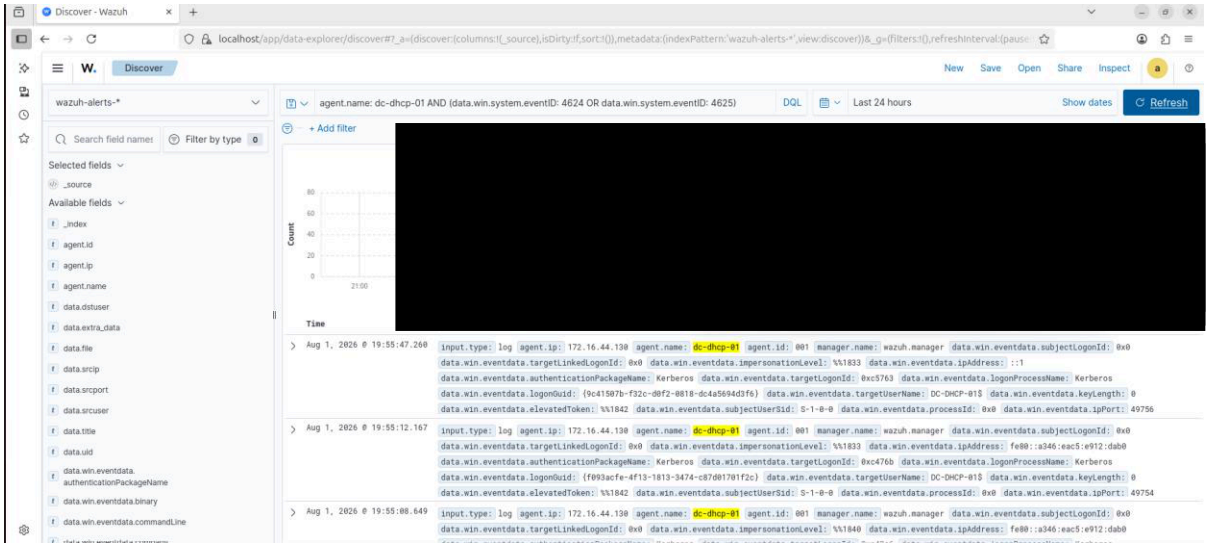


Figure 21 Domain Controller agent collecting no Windows telemetry

Figure 21 shows the Domain Controller agent is enrolled and active, its shared configuration contains no directives to collect Windows Event Log or Sysmon channels. The lateral movement and domain administrator logon therefore produced no events in the SIEM. The most sensitive host in the environment was effectively unmonitored at the point of compromise. This gap is addressed by adding event-collection directives during remediation, although live validation was later constrained as described in Section 7.

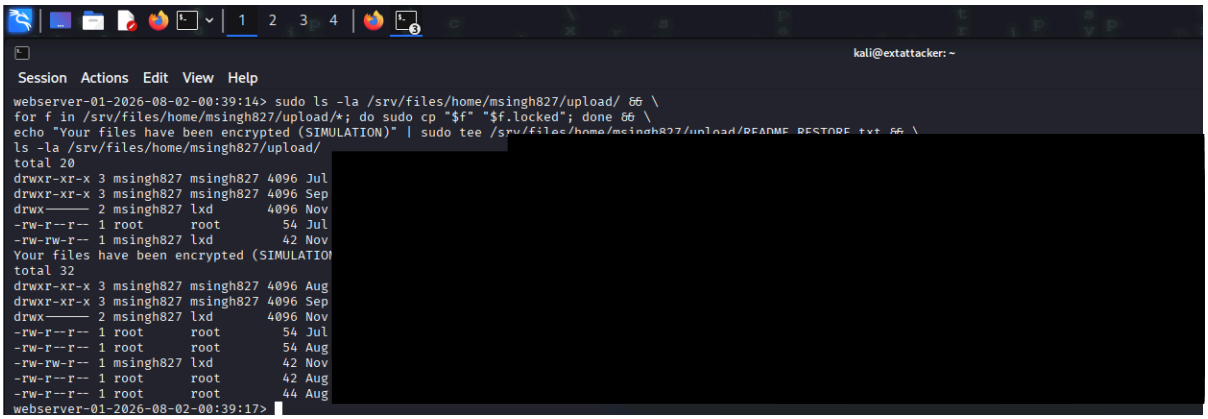


Figure 22 Empty Wazuh event query for the Domain Controller agent

Figure 22 shows querying the SIEM indexer for any alerts attributed to the Domain Controller agent across the exercise window returns no results, confirming that the domain administrator authentication and SMB activity produced zero detections. This empty result is the definitive evidence that the DC was unmonitored during the compromise. It distinguishes the absence of alerts from a suppressed or filtered result. Together with the preceding agent-configuration figure it provides the two-layer evidence needed to close the finding.

5.8 Impact

The chain concluded with a controlled ransomware simulation against the file share.

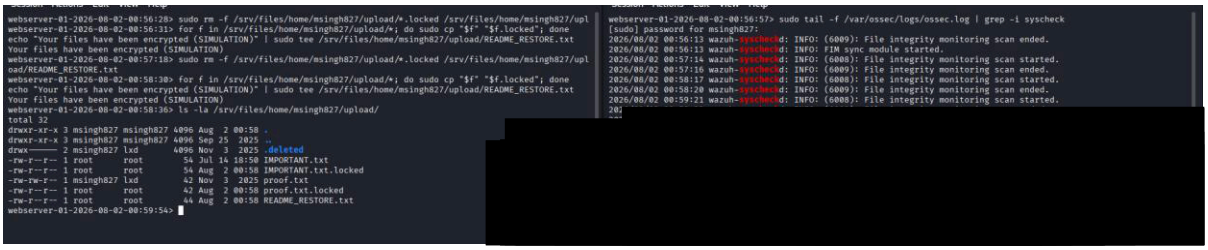


Figure 23 Ransomware simulation against the file share

Figure 23 shows A controlled routine copies documents on the file share to a locked extension and writes a ransom note, reproducing the observable impact of ransomware safely. No real malware or encryption was used, in keeping with the rules of engagement. The step corresponds to the Impact stage of the intrusion and reflects the outcome most feared by the target sector. In the baseline pass the burst of file changes produced no alert because the share was not yet monitored.

6 Detection Analysis and Findings

Applying the framework's three questions to each technique produced the baseline detection picture summarised in Table 3. The dominant pattern is that owned tooling was present but not configured to observe the attack surface, so most techniques passed with no prevention and no alert.

Table 3 Baseline detection outcomes per technique (P prevented, A alerted, L logged).

Technique	Expected detection	Prev.	Alert	Log
T1595 Scan	IDS scan signature	No	No	No
T1190 Exploit	Web application firewall	No	No	No

Technique	Expected detection	Prev.	Alert	Log
T1505.003 Web shell	File integrity monitoring	No	No	No
T1053.003 Cron	Endpoint monitoring	No	No	No
T1552.001 Credentials	Endpoint monitoring	No	No	No
T1021.004 SSH to SIEM	SSH auth rule	Partial	No	Host only
T1562.001 Impair defenses	Service-stop detection	No	No	No
T1021.002 SMB to DC	Windows logon + Sysmon	No	No	No
T1486 Impact	FIM mass-change rule	No	No	No

The analysis produced eleven findings, summarised with recommendations in Table 4. Two are critical: the domain administrator compromise through credential reuse, and the unauthenticated remote code execution achieved by chaining web-application vulnerabilities. The remainder concern missing telemetry, insecure credential storage, weak service configuration, and the resilience of the monitoring platform itself.

Table 4 Findings and recommendations.

ID	Finding	Severity	Recommendation
F-01	Domain administrator compromise via credential reuse	Critical	Rotate credentials, enforce unique local admin passwords, restrict spraying
F-02	Unauthenticated RCE via chained web vulnerabilities	Critical	Front the application with a WAF; patch or retire the vulnerable app
F-03	Credentials stored in files and disclosed in HTML	High	Remove secrets from files and comments; use a secrets store
F-04	No Windows telemetry reaching SIEM from the DC	High	Configure Windows Event Log and Sysmon collection on the DC
F-05	File server had no FIM on the attack surface	High	Extend FIM to the web root and file share
F-06	SIEM reachable over SSH from an untrusted segment	Medium	Restrict SSH by source; retain MFA
F-07	SIEM does not monitor its own authentication log	Medium	Ingest the host auth log; alert on SIEM logins

ID	Finding	Severity	Recommendation
F-08	SMB null session enabled on the DC	Medium	Disable anonymous SMB enumeration
F-09	Network IDS lacks scan detection	Medium	Add a threshold-based scan-detection rule
F-10	No pre-incident backup or recovery process	Medium	Establish and test regular backups
F-11	Monitoring lacks resilience to component restart	Low	Enable agent auto-reconnect and manager high availability

7 Remediation

Four control gaps were remediated during the exercise, each with a small, specific change that could be re-tested. A fifth item, telemetry on the Domain Controller, was configured but could not be fully validated live, for reasons documented below.

7.1 File Server Telemetry

File integrity monitoring was extended on the file server to cover the web root and the file share. Real-time notification was found not to be supported by the underlying share filesystem, so scheduled scanning at a short interval was used instead, which is a legitimate and reliable control.

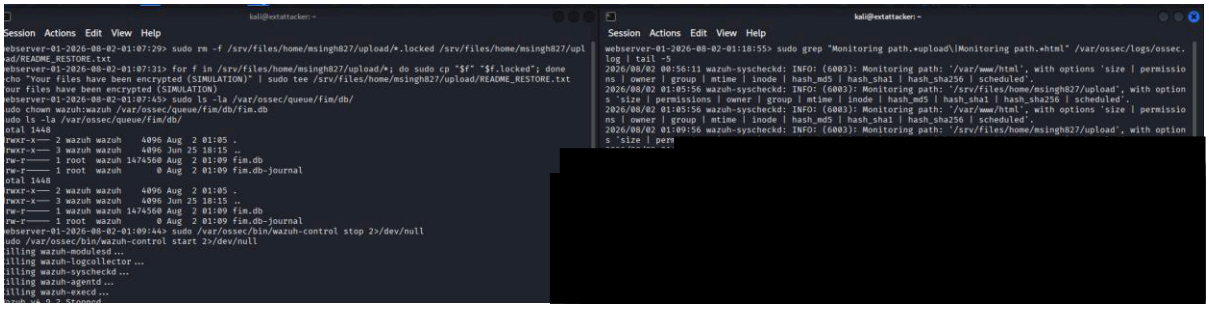


Figure 24 File integrity monitoring confirmed on the file server

Figure 24 shows the agent log and configuration on the file server confirm that file integrity monitoring now covers the web root and the file share, with full attribute and hash checking on a short scheduled interval. Scheduled scanning was selected because the share filesystem did not support real-time notification. This closes the endpoint-visibility gap on the file server that allowed the earlier activity to pass unseen. It applies directly to the paths used during the intrusion.

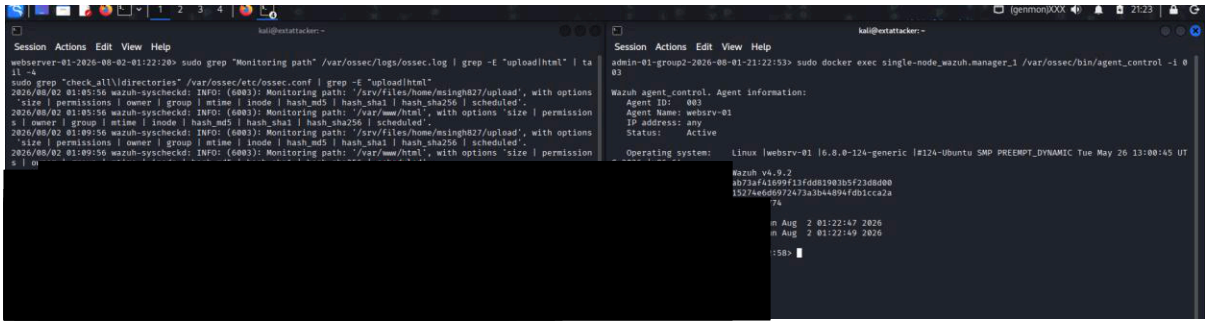


Figure 25 File server agent active with monitoring applied

Figure 25 shows the agent status and recent scan timestamps confirm the file server is active and reporting to the SIEM with the updated monitoring configuration in effect. This demonstrates that the remediation was applied to a live, connected endpoint rather than to configuration alone. It provides the baseline against which the validation re-run is measured. The platform is healthy and inspecting the intended paths.

7.2 Web Application Firewall

A reverse proxy running a web application firewall with the OWASP Core Rule Set, supplemented by custom command-injection rules, was placed in front of the previously unprotected application.

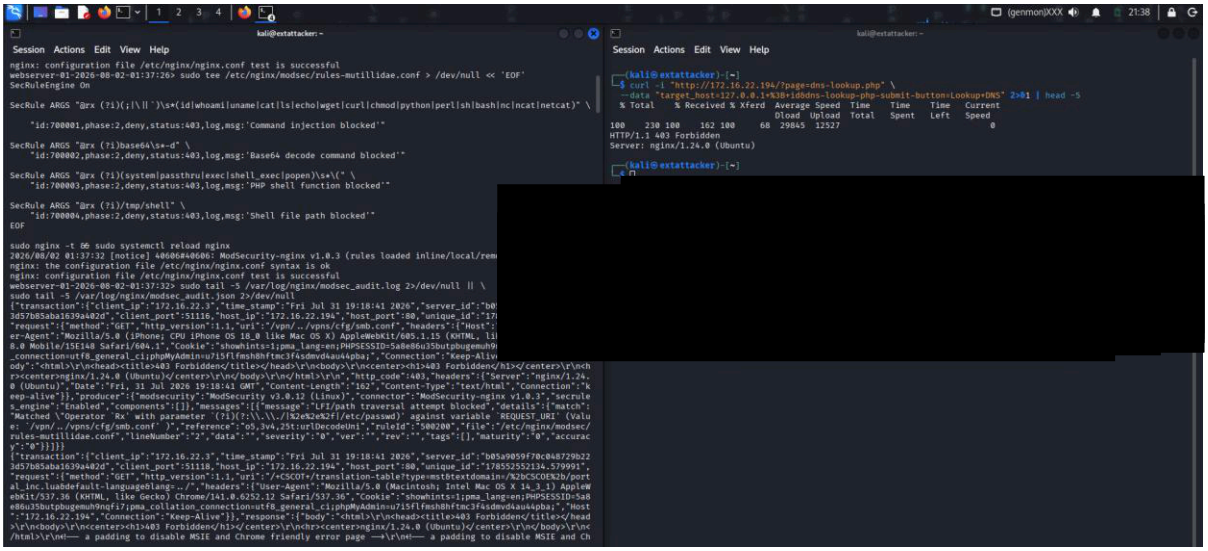


Figure 26 Web application firewall blocking the exploit

Figure 26 shows the same command-injection payload used to gain the initial foothold now returns a forbidden response after a reverse proxy with a web application firewall was placed in front of the application. The firewall inspects the request, matches a rule for the injected command pattern, and blocks it before it reaches the application. This closes the perimeter gap that previously allowed unrestricted exploitation. The control is verified using the exact request that succeeded during the baseline pass.

7.3 SIEM Hardening

The SIEM host was restricted so that SSH is accepted only from a management source, and a custom high-severity detection rule was added to alert on any successful login to the monitoring host.

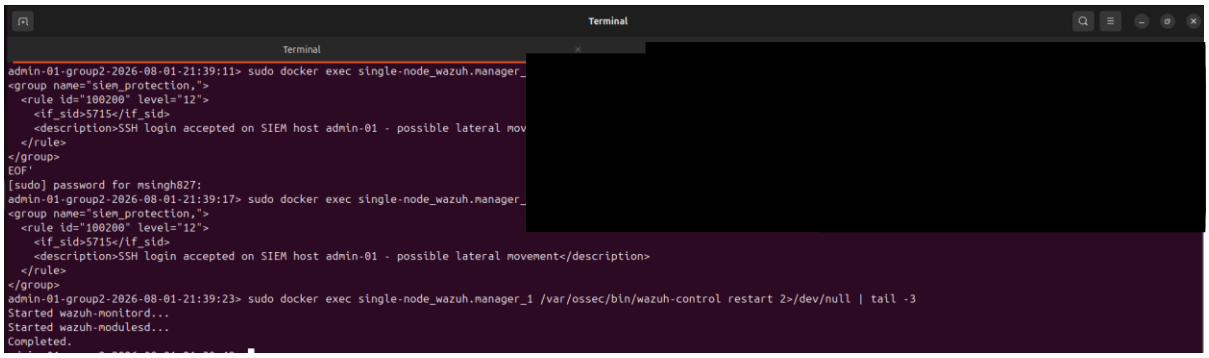


Figure 27 SIEM self-protection rule validated

Figure 27 A custom detection rule for authenticated access to the SIEM is validated against a representative log line, alongside a real blocked attempt from the attacker host. The rule engine confirms the rule fires at high severity with a description identifying possible lateral movement to the monitoring platform. Combined with a host firewall restriction, this closes the earlier blind spot around the SIEM. The environment now both limits exposure and raises a high-severity alert on any successful login to the monitoring host.

7.4 Backup and Recovery

Baseline backups of the web root and file share were created to a protected local path, establishing the recovery capability that the environment previously lacked.

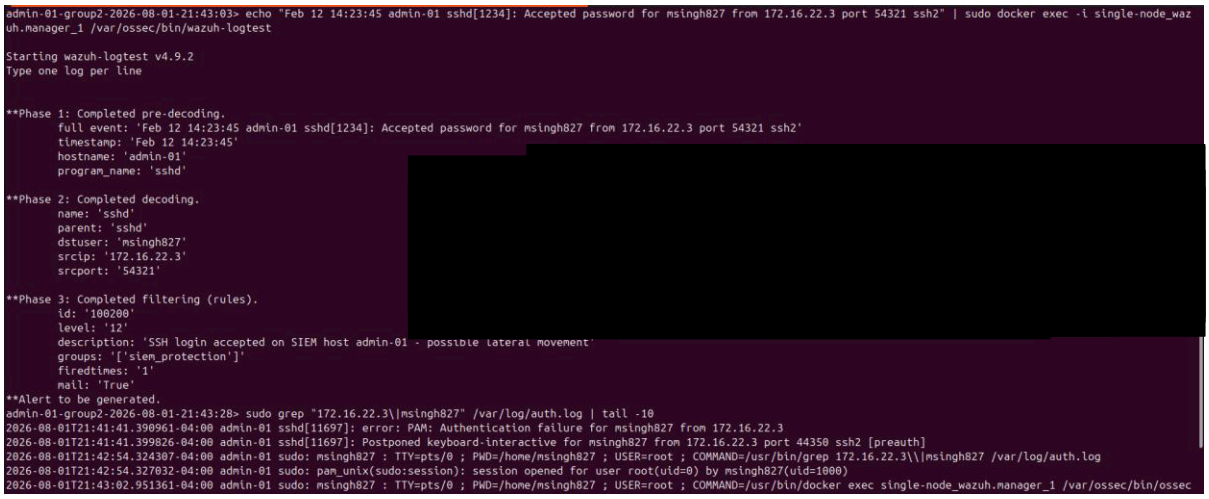


Figure 28 Baseline backup established

Figure 28 shows the evidence shows the creation of baseline backups of the web root and file share to a protected local path. Prior to this step no pre-incident backup existed, leaving recovery from a destructive event uncertain. The archive provides a known-good copy that can restore the share after the ransomware simulation. This establishes the recovery capability validated in the final phase.

7.5 Domain Controller Telemetry

Windows Event Log and Sysmon collection directives were added to the shared agent configuration for the Domain Controller. However, the Domain Controller agent had disconnected during the earlier defense-impairment step when the manager was restarted, and it could not be reliably reconnected within the exercise window because endpoint protection on the host blocked the remote service restart. Live validation of Domain Controller telemetry was therefore not achieved. This outcome is reported honestly and is itself

consistent with finding F-11 on monitoring resilience; the configuration change is in place for future collection.

8 Validation

The attack chain was re-run with the gaps closed to produce direct before-and-after evidence. Table 5 summarises the change in outcome for each stage.

Table 5 Before-and-after validation summary

Attack stage	Baseline (Pass 1)	After remediation (Pass 2)
Reconnaissance	Not detected	Not detected (rule pending)
Exploit web application	Succeeded, not blocked	Blocked with forbidden response
Web shell / file change	Undetected	File integrity alert raised
SSH to SIEM	Reached host, MFA blocked, no alert	Denied and high-severity rule fires
SMB to Domain Controller	Domain admin, no telemetry	Telemetry configured, not live-validated
Ransomware impact	Undetected, no recovery	Detectable and restored from backup



Figure 29 Exploit blocked on the validation re-run

Figure 29 shows the validation re-run of the initial-access attempt shows the previously successful payload again rejected with a forbidden response by the web application firewall. Repeating the exact baseline request confirms the control is effective rather than incidental. This provides direct before-and-after evidence that the initial-access gap is closed. The application is no longer exploitable through the request path used to gain the original foothold.



Figure 30 ModSecurity WAF setup and rule configuration

Figure 30 shows the configuration output shows the nginx reverse proxy loaded with ModSecurity rules covering command injection patterns, confirming the control is in place before the validation re-run. It identifies the specific rules targeting the semicolon-separated shell command pattern used during the

baseline exploitation. Showing both the configuration and the block response provides layered evidence that the WAF was correctly deployed rather than returning a coincidental error. This closes the initial-access gap with documented technical depth.

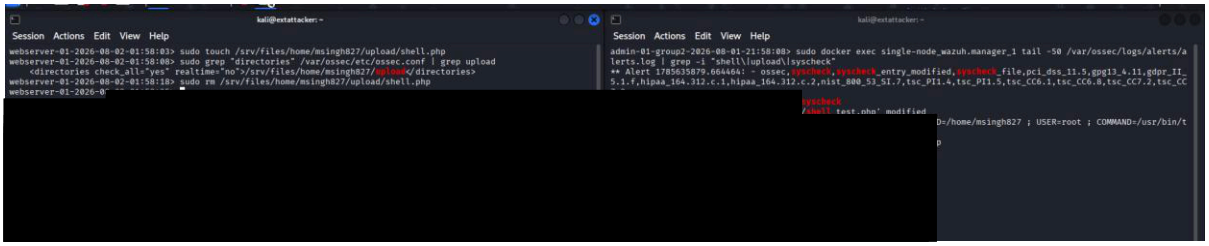


Figure 31 File change now detected by the SIEM

Figure 31 shows creating a file in the monitored share now generates a file integrity alert in the SIEM, the detection that was absent during the baseline pass. The alert records the modified path and the responsible event, confirming the monitoring pipeline functions end to end. This validates that a ransomware-style burst of file changes would now be surfaced to analysts. It is concrete evidence that the endpoint-visibility remediation is effective.

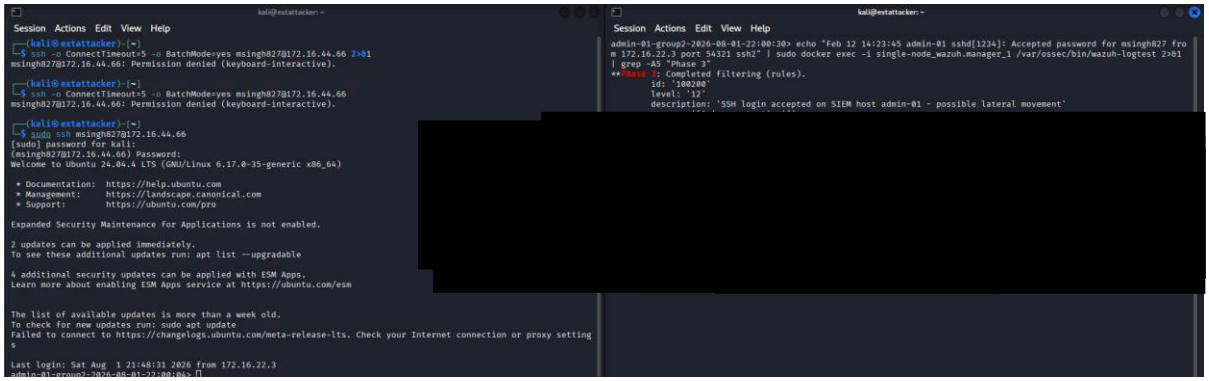


Figure 32 Lateral movement to the SIEM blocked and alertable

Figure 32 shows the repeated SSH attempt from the attacker host to the SIEM is denied, alongside confirmation that the self-protection rule would raise a high-severity alert on any successful login. Together these demonstrate both prevention and detection coverage for the monitoring host. The dual-role weakness identified during the baseline pass is therefore mitigated on both fronts. The SIEM is no longer a soft lateral-movement target without corresponding visibility.

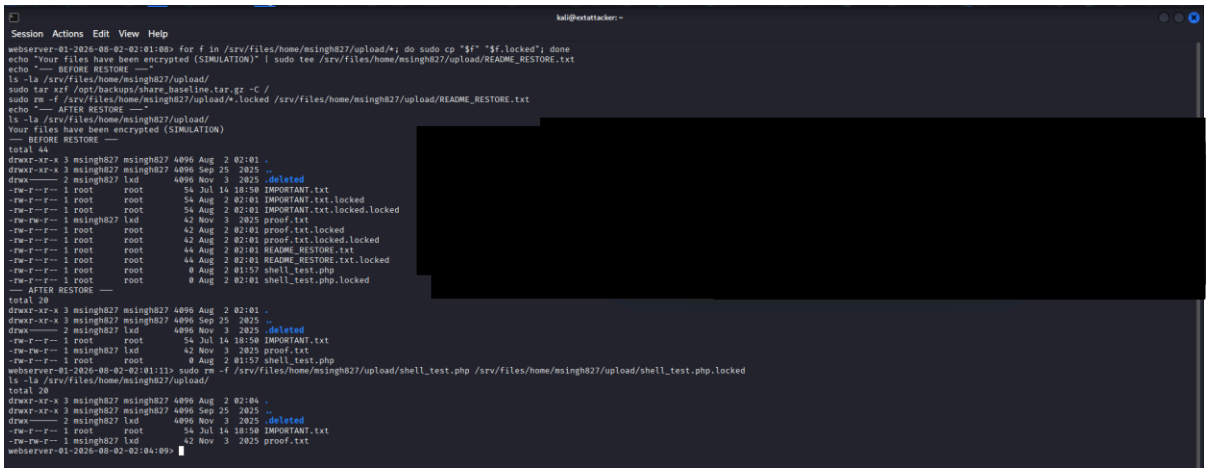


Figure 33 Ransomware impact reversed from backup

Figure 33 shows the ransomware simulation is re-run and then reversed by restoring from the baseline backup, after which the locked files are removed and the share returns to its clean state. The before-and-after listing confirms the destructive change was fully reversed using the recovery capability introduced during remediation. This validates that the organisation could recover from an encryption event without paying a ransom. It closes the recovery gap that left the baseline environment without a viable restore path.

9 Discussion and Lessons Learned

The exercise reinforced that the difference between owning security tooling and operating it is the essence of Maturity Level 1. Every host in scope carried an agent, an IDS watched the network, and a SIEM collected events, yet the baseline intrusion advanced from an exposed web application to full domain compromise without generating an actionable alert, because the tooling was not configured to observe the paths the attacker used. The most valuable output of the assessment was therefore not the successful attacks but the precise, evidenced account of where visibility was missing.

Two lessons stand out. First, the defense-impairment step demonstrated that impairing a single central component can cascade: restarting the manager disconnected every agent, and one endpoint did not recover, which both illustrates a genuine resilience weakness and cautions that destructive steps in a live exercise carry operational consequences that must be sequenced carefully. Second, container boundaries repeatedly limited host-based visibility, from the web root that host file integrity monitoring could not see, to the manager that could not read the host authentication log; monitoring architecture must account for where data actually resides rather than assuming a flat host view.

The assessment was conducted and documented honestly. Where a control behaved unexpectedly, such as multi-factor authentication blocking the SSH pivot, or where validation could not be completed, such as Domain Controller telemetry, the outcome was recorded as it occurred. This transparency is consistent with the purpose of a purple team, which is to give an organisation an accurate picture of its posture rather than a flattering one.

10 Conclusion

This purple team assessment executed a realistic, connected intrusion against a Maturity Level 1 healthcare environment, measured detection at every step, remediated four control gaps, and validated the improvement with a second pass. The baseline pass exposed critical exposure, including chained web-application exploitation and a domain administrator compromise through credential reuse, against a monitoring stack that produced almost no actionable signal. Remediation closed the initial-access, endpoint-visibility, SIEM-exposure, and recovery gaps, and validation confirmed that the exploit is now blocked, file changes are detected, lateral movement to the SIEM is prevented and alertable, and the ransomware impact can be reversed from backup. The organisation moved measurably from Maturity Level 1 toward a posture in which detection speed and recovery, the metrics that matter most to the healthcare sector, are demonstrably improved. The remaining items, particularly Domain Controller telemetry and monitoring resilience, define a clear and actionable path for continued maturation.

References

Center for Internet Security. (2025). *CIS critical security controls* (Version 8.1). https://www.cisecurity.org/controls/

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). *Computer security incident handling guide* (NIST Special Publication 800-61, Rev. 2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r2

Defensive Origins. (n.d.). *Atomic purple team framework and lifecycle*. GitHub. Retrieved from https://github.com/DefensiveOrigins/AtomicPurpleTeam

Joint Task Force Transformation Initiative. (2012). *Guide for conducting risk assessments* (NIST Special Publication 800-30, Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-30r1

MITRE Corporation. (2026). *MITRE ATT&CK: Enterprise matrix* (Version 19). https://attack.mitre.org/matrices/enterprise/

Open Information Security Foundation. (n.d.). *Suricata user guide*. https://suricata.readthedocs.io/

OpenVPN, Inc. (n.d.). *OpenVPN community resources and documentation*. https://openvpn.net/community-resources/

OWASP Foundation. (2021). *OWASP top 10: 2021*. https://owasp.org/Top10/

OWASP Foundation. (n.d.). *OWASP ModSecurity core rule set*. https://coreruleset.org/

SCYTHE. (n.d.). *Purple team exercise framework (PTEF)*. https://scythe.io/purple-team-exercise-framework

Wazuh, Inc. (n.d.). *Wazuh documentation*. https://documentation.wazuh.com/