



Prifysgol Cymru

Y Drindod Dewi Sant

University of Wales

Trinity Saint David

Submitted by

Shamsuddin Mahmud (2432880)

Supervisor: Dr. Mohamed Ghanem

Project submitted as part of the requirements for the
award of MSc

TASK - 1

ALPHV/BlackCat Ransomware Attack on Change Healthcare

Table of Contents

Introduction	3
Attack Timeline	3
MITRE ATT&CK TTPs	4
Root Cause Analysis.....	6
Strategic Recommendations	7
Conclusion	9
References.....	9

Introduction

A significant ransomware incident occurred at Change Healthcare, a subsidiary of UnitedHealth Group that processes approximately one-third of U.S. healthcare transactions, in February 2024, resulting in nationwide impacts. The attackers were reportedly part of the ALPHV/BlackCat ransomware group, a well-known ransomware-as-a-service (RaaS) gang. The attack is believed to be the most significant cyberattack on the U.S. healthcare system to date. The American Hospital Association Leadership had stated it as "the most significant and consequential incident of its kind against the U.S. healthcare system in history" [1]. In breaching the company, the attackers obtained a significant trove of sensitive data (reported to range as high as six terabytes of information, including patient records and Social Security numbers). They ultimately locked up life-saving equipment, devastating healthcare services in all 50 states [1]. UnitedHealth Group ultimately paid USD 22 million, although only a small portion of the stolen information was returned [2]. In total, the breach compromised personal health information for approximately 190 million individuals, making it the most significant healthcare data breach ever recorded in the United States [3]. This article breaks down the timeline, TTPs of the attack as well as their features and how they can be countered, but also we will uncover the root causes and strategic mitigations to avoid similar situations in future.

Attack Timeline

- **Early to Mid-February 2024 - Initial Breach:** On 11-12 February 2024, attackers affiliated with BlackCat gained initial access to Change Healthcare's network. According to congressional testimony, the hackers used compromised user credentials to remotely log into a Citrix remote access portal that lacked multifactor Authentication (MFA) [2]. The intruders then spent the next nine days inside the network, moving laterally and stealing data, remaining undetected by the organisation's security monitoring [3].
- **Feb. 21, 2024 - Ransomware Detected:** In early February, the BlackCat ransomware was executed on Change Healthcare's servers, encrypting files across the enterprise [2][3]. Change Healthcare's IT staff discovered the ongoing attack that day and began shutting down systems and taking equipment offline to halt the further spread of the malware [1]. Such a sudden closure of core platforms effectively suspended the majority of health service provision not only in the Hagaz sub-zone, but also nationwide.
- **Late February 2024 - Public Disclosure and Threats:** By Feb. 26, the **BlackCat gang claimed responsibility** for the attack via their leak site, threatening to publish stolen data if their demands were not met [1]. On Feb. 29, Change Healthcare officially confirmed BlackCat's responsibility. It revealed that approximately **6 TB of data** had been stolen, including patient Social Security numbers, medical records, and even records involving military personnel [1].

- **Early March 2024 - Ransom Payment and Mitigation Efforts:** Facing mounting pressure, UnitedHealth Group (Change Healthcare's parent) authorised a ransom payment of 22 million in Bitcoin to BlackCat in early March 2024 [2]. Even though the ransom was paid, the hackers still did not return the data or decryption keys, a sobering reminder that paying a ransom is no guarantee of recovery. At the same time, UnitedHealth's Optum division quickly created an emergency financial assistance programme for hospitals and clinics that had been affected by the outage, providing them with zero-interest loans to ease cash-flow problems while their claims systems were inoperable.
- **April 2024 - Secondary Extortion Attempt:** In April 2024, a second ransomware group, known as "RansomHub," emerged, claiming to have obtained a copy of Change Healthcare's stolen data (likely from a disgruntled BlackCat affiliate) and attempted further extortion [1]. RansomHub released sample data and threatened to auction the entire haul to the highest bidder, highlighting how data can remain exposed even after a ransomware attack has been contained.
- **Mid-to-late 2024 - Recovery and Impact Assessment:** Recovery of services. At the end of May, Change Healthcare began recovering over a period of weeks and months; core clearinghouse operations were reported as fully operational by late 2024. However, the downstream effects persisted for months, with many providers still struggling to clear their billing backlogs and facing financial difficulties. By July 2024, Change Healthcare began issuing breach notifications, initially estimating that around 100 million individuals were affected, a number that continued to climb as data analysis progressed [3]. By early 2025, federal filings indicated that approximately **190 million patients and consumers** had their data compromised - more than half the U.S. population - dwarfing any previous healthcare breach [3]. In the aftermath, the incident also led to numerous lawsuits and multiple state and federal investigations, reflecting the widespread fallout of this attack.

MITRE ATT&CK TTPs

The ALPHV/BlackCat campaign against Change Healthcare, for example, illustrated a variety of TTPs mapped to the MITRE ATT&CK framework. Ten, and possibly more separate TTPs from the entry point to reach, that can be observed or derived from the attack:

- **Initial Access - Valid Accounts & External Remote Services (T1078, T1133):** The attackers used compromised user credentials to log in to an unsecured Citrix remote desktop gateway that did not have multifactor authentication [2]. Leveraging an external service that was not protected by a valid account, the attackers gained access to our defences and established an initial foothold.
- **Execution - Command and Scripting Interpreter (T1059.003):** Once inside, BlackCat likely executed commands and scripts to control systems and deploy malware. BlackCat ransomware is known to run via the Windows command shell (cmd.exe) with specific parameters to initiate encryption [4]. The use of built-in command-line tools helps attackers to blend into everyday operations.

- **Persistence - Modify Registry (T1112):** BlackCat can establish persistence by modifying system registry settings. For instance, the malware has been observed adding or modifying registry keys related to network and file-sharing parameters to facilitate lateral movement and mass encryption across shared resources [4]. By persisting through registry changes, the attackers ensure their malware or backdoors restart even after reboots.
- **Privilege Escalation - Abuse Elevation Control Mechanism (T1548.002):** To fully compromise high-value servers, the attackers likely escalated privileges. BlackCat ransomware can **bypass User Account Control (UAC)** on Windows, allowing it to run with administrative rights [4]. With elevated privileges, the malware can turn off security tools and access more sensitive data.
- **Defence Evasion - Indicator Removal on Host (T1070.001):** Clearing logs and system traces is a technique employed by BlackCat to evade detection. The malware can wipe Windows event logs using utilities like wevtutil.exe, effectively erasing evidence of its activities on infected hosts [4]. This hinders incident responders' ability to reconstruct what happened.
- **Discovery - Network Share and Account Discovery (T1135, T1087.002):** During the dwell time, the intruders performed internal reconnaissance. BlackCat is known to enumerate network shares and user accounts, for example, using net use commands to map out the domain environment [4]. In this attack, the affiliate likely identified key file servers, databases, and backups to maximise the damage from encryption.
- **Lateral Movement - Lateral Tool Transfer (T1570):** To spread across the enterprise, the attackers employed lateral movement techniques. BlackCat affiliates commonly use tools like PsExec (a Windows remote execution utility) to propagate ransomware to multiple servers in parallel [4]. Moving laterally ensured that almost all of Change Healthcare's systems, including databases and applications, were eventually infected and encrypted on Feb. 21.
- **Collection/Exfiltration - Data Staging and Theft:** Before triggering encryption, the threat actors stole a large volume of data (personal and health records totalling several terabytes) as part of a double-extortion strategy [1]. They likely compressed and transmitted these files out of the network to servers under their control, thereby obtaining leverage for extortion.
- **Impact - Data Encrypted for Impact (T1486):** Finally, the attackers executed the ransomware payload to encrypt data on thousands of systems nearly simultaneously [2]. By encrypting databases, application servers, and workstations, they effectively paralysed Change Healthcare's operations. BlackCat's malware uses robust encryption (combining RSA and AES) implemented with multi-threading for speed, making recovery without the decryption keys virtually impossible [4].
- **Impact - Inhibit System Recovery (T1490):** In tandem with encryption, BlackCat attempts to inhibit recovery by deleting backups and shadow copies on Windows systems [4]. The ransomware issues commands like vssadmin.exe delete shadows /all /quiet and

may alter boot configurations to turn off restoration. This ensures that the victim cannot easily restore data from local backups, increasing pressure to pay the ransom.

All in all, the Change Healthcare attack also represents a full-spectrum attack that combines advanced malware capabilities with basic misconfigurations (in this case, a lack of MFA). Credential theft, invisible lateral movement, data exfiltration, and destructive encryption all match BlackCat's reported tactics, according to the threat intelligence [4].

Root Cause Analysis

Compromised Access Controls: The root cause of this breach can be traced to a failure in identity and access management. Change Healthcare's Citrix remote access server allowed logins with only a username and password, without requiring multifactor authentication [2]. Attackers exploited this outdated configuration, utilising valid credentials (likely obtained through phishing or purchased on underground markets). By taking advantage of the absence of MFA - or two-factor authentication, a standard security feature - the attackers succeeded in entering as an authenticated user, bypassing any company perimeter defences. In short, one pilfered password gave the invaders the keys to the entire network.

Network Security Gaps: The attackers remained undetected for nine days, indicating insufficient monitoring and a flat network structure. Poor internal segmentation and weak anomaly detection allowed the intruders to move laterally into core systems without triggering alarms.

Data Protection Failures: Another underlying issue was the lack of adequate data-level protections. The exfiltrated data was highly sensitive (including PHI, such as insurance IDs, diagnoses, and SSNs). However, much of it was stored in an unencrypted form within the network. Once the intruders gained access, they could collect vast troves of patient data. This points to inadequate encryption of data at rest and insufficient data loss prevention controls. Moreover, backup systems, if they existed, were not sufficiently protected (e.g. not segregated offline); the attackers were able to delete or encrypt backups, as evidenced by the prolonged outage. This indicates a lapse in disaster recovery preparedness.

Third-Party Risk Oversight: The magnitude of this breach also reflects inadequate oversight of a critical vendor. Neither regulators nor client hospitals ensured that basic cybersecurity controls, such as multifactor Authentication (MFA) and network segmentation, were in place at Change Healthcare. In retrospect, risk assessments and compliance audits proved to be insufficient. A provider handling such a large share of healthcare transactions should have been held to higher security standards, given that its compromise had **systemic consequences**.

Strategic Recommendations

To prevent a recurrence of an attack of this scale, Change Healthcare (and similar organisations) must implement comprehensive improvements. The following strategic recommendations span architectural design, technical controls, and monitoring enhancements:

Architectural Improvements

- **Zero Trust Networking:** Adopt a zero-trust architecture where no user or system is inherently trusted. This involves **strict network segmentation** - isolating critical systems (e.g. claims processing and database servers) in protected enclaves - and requiring re-authentication and authorisation for lateral movements. By compartmentalising the network, even if one node is breached, the blast radius is minimised.
- **Resilient Infrastructure Design:** Build redundancy and fail-safes into core systems. For example, maintain **offline, air-gapped backups** of critical databases and files, and regularly test restoration from these backups. Ensure that backup credentials and systems are segregated from the main network domain. Thus, even if production systems are encrypted by ransomware, operations can be restored from clean backups.
- **Secure or Decommission Legacy Systems:** Identify any legacy or unsupported systems (such as older servers or applications lacking modern security features) and either retire them or isolate and harden them. Any remote access portals that cannot enforce MFA or other controls must be upgraded or heavily restricted. Architectural decisions should prioritise up-to-date platforms that offer built-in security and easier patch management.

Technical Controls

- **Strong authentication everywhere:** Enforce **multifactor authentication** on all user accounts, especially for remote access and privileged logins. Had MFA been in place on the Citrix gateway, the initial breach likely would have been thwarted. This measure is strongly recommended by cybersecurity authorities [5].
- **Endpoint and Network Security:** Deploy advanced **Endpoint Detection and Response (EDR)** tools on servers and workstations to detect ransomware behaviours (such as rapid file encryption or mass deletion of shadow copies) and halt them in real time. Likewise, implement network-level defences, including intrusion detection/prevention systems (IDS/IPS), to monitor lateral traffic. Consider using application allow-listing to prevent unauthorised programs (like unknown binaries) from executing on critical servers [5].
- **Continuous Vulnerability Management:** Institute a rigorous patch management programme to promptly address known vulnerabilities in operating systems, remote access software, and other applications. Ransomware groups often exploit unpatched systems as entry points [5]. By staying current with security updates and using automated scanning, the organisation reduces its attack surface. Additionally, perform periodic penetration tests to identify and remediate any configuration weaknesses before attackers find them.
- **Data Encryption and DLP:** Encrypt sensitive data at rest within databases and file storage, so that even if data is exfiltrated, it remains protected. Implement Data Loss

Prevention (DLP) solutions to flag and block unusual bulk data transfers, especially involving protected health information. For instance, significant exports of patient records should trigger alerts or require approval. These measures can limit attackers' ability to steal massive datasets undetected.

- **Incident Response Preparedness:** Develop and regularly practice an **incident response and disaster recovery plan** tailored to ransomware scenarios. This plan should define clear roles, communication protocols (including those with law enforcement, regulators, partners, and media), and technical steps for isolation, containment, and recovery. Conduct regular tabletop exercises or simulations of a ransomware attack to ensure the team can respond swiftly and effectively to such an event.

Monitoring and Detection

- **24/7 Security Operations Centre:** Establish a dedicated Security Operations Centre (SOC) with around-the-clock monitoring of logs and alerts. Aggregate logs into a central SIEM (Security Information and Event Management) system to enable real-time analysis of suspicious patterns (e.g., an account logging into numerous machines or large volumes of data being transferred out). In this incident, earlier detection of the attackers' lateral movement or data exfiltration could have drastically limited the damage.
- **User and Entity Behaviour Analytics:** Implement **behavioural analytics** to detect anomalies in user and system activity. For example, an employee account initiating an atypical remote session or accessing large data repositories at odd hours should raise immediate flags. Machine-learning-based tools can establish a baseline of normal behaviour and alert on deviations that may indicate an intruder using compromised credentials.
- **Threat Intelligence and Hunting:** Leverage threat intelligence feeds (including up-to-date indicators of compromise for BlackCat/ALPHV) to update defensive systems. Security teams should proactively hunt for signs of known TTPs - such as creation of suspicious registry keys, use of PsExec, or presence of BlackCat-specific binaries - within the environment. Being threat-informed enables defenders to catch adversaries in the reconnaissance or early exploitation phase.
- **Regular Security Audits:** Conduct periodic audits of security controls to ensure they are covering all critical systems and functioning as intended. Verify that all servers, cloud workloads, and endpoints are generating logs and that alert rules are tuned to minimise false negatives. Use findings from post-incident analyses to strengthen monitoring. Continuous improvement of detection capabilities will help catch sophisticated threats before they escalate and become more severe.

By implementing the above architectural, technical, and monitoring measures, Change Healthcare can significantly bolster its cybersecurity posture. These steps align with expert guidance from federal agencies for defending against ransomware groups, such as BlackCat [5]. Additionally, the organisation should foster a security-first culture through ongoing staff training (to resist phishing and social engineering) and executive oversight of cyber risk.

Conclusion

The BlackCat ransomware attack on Change Healthcare in February 2024 was a wake-up call for the healthcare industry. The use of single-factor authentication, inadequate internal controls and layering of security allowed a cybercriminal collective to wreak havoc in an unprecedented manner, effectively taking a significant health care system hostage. The aftermath, which consisted of approximately \$2.5 billion in response costs and the compromise of 190 million health records, serves to reinforce the notion that cybersecurity is more than just an IT problem; it is also an issue of patient safety and national security.

The incident has, in hindsight, shown that even the best-funded can succumb to basic security oversights, and the knock-on effect of such an attack could pose a threat to healthcare itself. However, it is helpful to learn the lessons. Through the rigorous adoption of practices such as zero-trust architecture, strong authentication, continuous monitoring, and well-rehearsed incident response protocols, organisations can significantly minimise the chances of suffering a ransomware attack. The Change Healthcare incident has sparked industry-wide demands for more accountability and stronger security standards among healthcare IT vendors. Proactively and comprehensively implementing an effective cybersecurity plan, while also conforming to frameworks such as MITRE ATT&CK for threat mapping, will become imperative in protecting critical healthcare infrastructure from the growing spectrum of cyber threats. As shown, the price of doing nothing significantly outweighs the amount required to ensure the operational integrity of these critical systems.

References

- [1] Hyperproof Team, "Understanding the Change Healthcare Breach and Its Impact on Security Compliance," *Hyperproof Blog*, Dec. 7, 2025. [Online]. Available: <https://hyperproof.io/resource/understanding-the-change-healthcare-breach/>
- [2] J. Gregory, "Change Healthcare discloses USD 22M ransomware payment," *IBM News*, May 24 2024. [Online]. Available: <https://www.ibm.com/think/news/change-healthcare-22-million-ransomware-payment>
- [3] S. Alder, "Nebraska AG's Lawsuit Against Change Healthcare Survives Motion to Dismiss," *HIPAA Journal*, Nov. 17 2025. [Online]. Available: <https://www.hipaajournal.com/change-healthcare-responding-to-cyberattack/>
- [4] MITRE ATT&CK, "BlackCat Ransomware (S1068) - Enterprise Techniques," MITRE, Last modified Oct. 21, 2025. [Online]. Available: <https://attack.mitre.org/software/S1068/>
- [5] FBI and CISA, "#StopRansomware: ALPHV/BlackCat Ransomware Advisory (AA23-353A)," Joint Cybersecurity Advisory, Dec. 19 2023.

TASK – 2

Table of Contents

Abstract.....	11
1. Introduction.....	11
2. Rules of Engagement (RoE)	11
3. Threat Emulation Plan and MITRE ATT&CK Mapping	12
4. Lab Environment Design (Hybrid Architecture)	13
5. Campaign Execution and Evidence	14
6. Detection Gaps and Blue Team Findings (Analysis).....	15
7. Strategic Recommendations and Control Enhancements	16
8. Cost-Benefit Analysis (Business Case).....	16
9. Conclusion	17
Appendix A. Evidence	18
References.....	35

Abstract

In this paper, we describe Operation IP Shield, a semi-automated red team campaign that models an attack attempt to exfiltrate placeholder pharmaceutical research data from an on-premises Active Directory environment, utilising a VPC-hosted portal and CI/CD process. Operations were scheduled and documented using the MITRE ATT&CK framework [1], which has subsequently been incorporated into the Cyber Kill Chain [11]. The evidence demonstrates how quickly reconnaissance and service fingerprinting can uncover an exposed attack surface, particularly under the right conditions, such as when server-side validation and secure SDLC controls are weak. It also illustrates how application-layer request tampering may occur in a case-specific scenario determined by this exposure. The report examines detection possibilities and common visibility blind spots in hybrid telemetry, as well as recommended remediation. A cost-benefit perspective is also employed to support informed executive decision-making.

1. Introduction

Pharmaceutical organisations are high-value targets because commercially sensitive intellectual property and regulated data are concentrated in a small number of systems and workflows. Sophisticated adversaries typically operate as campaigns, combining reconnaissance, credential abuse, and low-noise lateral movement, so that compromise can persist long enough to identify and collect the most valuable datasets. For this reason, assurance requires more than vulnerability scanning; it necessitates adversary emulation, which tests identity, cloud workloads, and software supply chains under realistic constraints.

Operation IP Shield was commissioned to emulate an APT-style intrusion targeting a specific sample of dummy research data, while minimising operational impact and ensuring that all activities remained within explicit authorisation. CPC's environment was described as hybrid: an on-premises Active Directory [10] network supporting R&D, an AWS VPC [8] hosting a clinical trial data portal, and a CI/CD pipeline supporting application development. These three domains, identity, cloud workloads, and software supply chain, commonly intersect during real-world compromises. For example, an exposed portal can enable credential reuse; a CI/CD pipeline can be abused to modify artefacts; and AD misconfigurations can accelerate privilege escalation and broad access.

The engagement was structured to be repeatable and defensible by mapping objectives and observed activities to MITRE ATT&CK [1] and summarising progression using the Cyber Kill Chain [11]. These frameworks enable CPC to translate technical evidence into concrete detection and control requirements, and to measure improvements over time.

2. Rules of Engagement (RoE)

A Rules of Engagement (RoE) document defines the legal, ethical, and operational constraints for a red team activity. For Operation IP Shield, the RoE was designed to preserve safety, prevent business disruption, and ensure results were audit-ready.

Scope and boundaries:

- In-scope systems: a lab-representative on-premises AD domain, an AWS VPC [8]

representative portal host, and a CI/CD demonstration environment supporting a sample application workflow.

- Data handling: only dummy research data was used; no real patient, employee, or proprietary datasets were accessed or exported.
- Allowed techniques: reconnaissance, enumeration, controlled vulnerability verification, and limited exploitation-style demonstrations strictly within the lab.
- Prohibited techniques: denial-of-service, destructive payloads, ransomware simulation, and any persistence beyond the approved test window.

Communications and approvals:

- A single engagement manager acted as the CPC point of contact, authorised scope changes, and coordinated the “stop” process.
- A safety stop procedure was defined (immediate cessation upon observed instability, unexpected system impact, or scope ambiguity).
- All artefacts (screenshots, logs, findings) were stored securely and provided only to CPC stakeholders.

Success criteria:

- Primary objective: demonstrate a plausible path to access a defined dummy research sample without causing disruption, and document evidence sufficient for blue team validation.
- Secondary objective: quantify where detection controls succeeded or failed and recommend prioritised improvements.

3. Threat Emulation Plan and MITRE ATT&CK Mapping

The threat emulation plan translated CPC’s mission objective (“protect R&D IP from a sophisticated competitor or nation-state”) into a phased campaign that can be measured and evaluated. The campaign assumes an adversary profile with moderate resources, strong operational security, and a focus on long-term access rather than immediate disruption. The plan prioritised stealth, credential-centric access, and low-noise discovery patterns, consistent with documented APT tradecraft.

Operational objectives were mapped to MITRE ATT&CK [1] to ensure traceability. High-level mapping is presented in Table 2; detailed technique references were used internally for test design and for recommending detection content (analytics and alerts). The plan explicitly included decision points where the campaign could shift from direct exploitation to identity-based abuse depending on observed defences and configuration posture.

Campaign Objective	Kill Chain Phase	ATT&CK Tactic	Example Techniques (non-exhaustive)
Enumerate exposed services and entry points.	Reconnaissance	Reconnaissance / Discovery	Active Scanning; Network Service Discovery; Service Version Discovery

Validate web application attack surface.	Delivery/Exploitation	Initial Access	Exploit Public-Facing Application; Valid Accounts (if available)
Assess identity boundary and AD exposure.	Installation/Command & Control	Discovery / Lateral Movement	Remote Services; SMB/LDAP Discovery; Kerberos/AD Enumeration
Demonstrate controlled data access and exfil simulation	Actions on Objectives	Collection / Exfiltration	Data from Information Repositories; Exfiltration Over Web Service (simulated)
Evaluate detection and response readiness.	Across phases	Defence Evasion / Detection	Indicator Removal on Host (not executed); Obfuscated/Compressed Information (reviewed)

4. Lab Environment Design

A representative lab environment was assembled to reflect CPC's hybrid composition while avoiding production risk. The goal was to reproduce the key trust boundaries (identity, cloud workload segmentation, and CI/CD permissions) so that attacker paths and defensive telemetry could be evaluated using realistic data flows and control-plane behaviour.

4.1 On-premises Active Directory [10] (R&D)

A lab AD domain controller and Windows hosts were configured to expose typical AD services (Kerberos, LDAP, SMB, RPC) within a segmented network. This enabled identity-focused discovery and assessment of common misconfiguration risks without interacting with production systems.

4.2 AWS VPC [8] (Clinical Trial Portal)

A representative web portal host was placed in an AWS-like segmented network model with security-group-style filtering and log sources analogous to cloud deployments. The design supported the evaluation of exposed web services, access controls, and cloud logging assumptions.

4.3 CI/CD Pipeline (Software Delivery)

A small CI/CD demonstration workflow was used to evaluate governance controls, permissions, secret handling, and artefact integrity rather than to pursue destructive pipeline abuse. Fig 17 evidences the lab CI/CD target context used during the assessment.

5. Campaign Execution and Evidence

This section summarises the campaign phases executed in the lab and links each phase to evidence screenshots. The screenshots are included to allow independent validation by the blue team and to support reproducing defensive detections without exposing sensitive organisational data.

5.1 Reconnaissance and Reachability Verification

The campaign began with basic connectivity checks to confirm routing between the assessment workstation and target segments. Figure 1 and 27 show ICMP reachability checks used to verify that the Kali host could communicate with the target web host and the AD host within the lab network. Such verification is a prerequisite for interpreting subsequent scan results (for example, distinguishing filtered ports from unreachable hosts).

5.2 Service Discovery and Fingerprinting Port and service discovery was conducted using Nmap [12] and Metasploit [7] `db_nmap` to inventory reachable services. Figures 2, 3, 19, 29, and 30 capture baseline port discovery and OS/service fingerprinting across the web host and the Windows/AD host. Fig. 4-9 provides a deeper version of probing and script-assisted enumeration to obtain actionable fingerprints, while Figs. 31-32 capture additional outputs that reveal domain-relevant metadata. These results established the attack surface that would realistically be available to an external or pivoted attacker and defined the telemetry CPC should expect to see during real reconnaissance activity.

From a defensive perspective, the value of this phase is twofold: (i) it identifies externally reachable attack surface (services that should be minimised and hardened), and (ii) it provides measurable opportunities for detection engineering (e.g., alerts on abnormal scanning, anomalous LDAP queries, or repeated failed service probes).

5.3 Controlled Vulnerability Verification and Exploitation-Style Demonstrations. Where fingerprints suggested that a known vulnerability might exist, controlled verification was performed in the lab to confirm the potential impact and to collect evidence for remediation. Fig 10-16 documents the workflow used to select a candidate exploit module, validate required parameters, review compatible payload families, and confirm an authorised test session. The demonstration was limited to the minimum actions necessary to demonstrate impact and did not attempt to persist, destroy data, or expand privileges. The security takeaway is that exposed and weakly governed services can provide rapid footholds that bypass later-stage controls unless compensating controls (patch governance, hardening, segmentation, and detection) are in place.

5.4 Web Application Proxying and Request Manipulation (DVWA [5] Demonstration) To model portal-style risks and validate application-layer controls, Burp Suite [6] was used as an intercepting proxy against DVWA [5] (an intentionally vulnerable application) in a controlled lab. Figure 20 shows the proxy in a baseline state. Fig 21: Evidence authentication behaviour. Fig 22-26 captures a representative HTTP request, inspection of body parameters, and controlled request modification to demonstrate how inadequate server-side validation can enable tampering and injection-style attacks. In production, equivalent risks are mitigated through secure coding (parameterised queries, robust input validation), WAF rules, and CI/CD security testing.

5.5 DNS and Domain Context Checks

Finally, DNS-oriented enumeration and domain context extraction were performed to evaluate whether naming services and directory metadata might expose practical attacker intelligence. Fig 33 shows the DNS test output used to verify exposure signals. These checks inform both hardening (restricting unnecessary information disclosure) and monitoring (detecting anomalous name resolution patterns).

6. Detection Gaps and Blue Team Findings

Operation IP Shield generated the most value when technical evidence could be effectively translated into control and detection requirements. The observations below summarise where hybrid organisations commonly lose visibility and why those gaps matter for an IP-theft threat model.

6.1 Inconsistent Telemetry Across Layers

Reconnaissance and request-tampering activity can be low-noise, yet it is detectable when telemetry is centralised. Without a SIEM and consistent log onboarding, scan bursts, anomalous authentication, and suspicious web parameters are often fragmented across systems, failing to trigger investigations. CPC should standardise the collection of endpoint, AD, application, and cloud control-plane logs and correlate them to ATT&CK-aligned analytics [1], [3].

6.2 Identity as the Primary Control Plane

The AD surface mapped in Fig 29-32 illustrates that identity services (Kerberos/LDAP/SMB/RPC) are both high value and highly observable. If privileged accounts lack MFA and tiering, a single credential can enable rapid lateral movement. Detection should focus on abnormal directory enumeration, unusual privileged logons, and high-risk SMB session patterns, combined with strong preventative controls (PAM, least privilege, segmentation) [3], [10].

6.3 Application-Layer Weaknesses as an Entry Point

The Burp Suite [6] workflow against DVWA [5] demonstrates that application flaws can bypass perimeter controls. In production, secure coding practices (including input validation and parameterised queries), CI/CD security testing, and WAF telemetry help mitigate this risk [4], [9]. Detection should include monitoring for anomalous parameters, repeated errors, and authentication anomalies at the application layer.

6.4 Change and Patch Governance

Exploitation-style validation reinforces that patch latency and unnecessary service exposure materially increase risk. A disciplined vulnerability management program, asset inventory, patch SLAs, configuration baselines, and exception governance reduce the feasibility of many initial access paths and shrink the window in which an attacker can act.

7. Strategic Recommendations and Control Enhancements

Recommendations are prioritised by risk reduction, feasibility, and cross-domain impact (on-premises, cloud, and CI/CD). They are framed as “control objectives” so that CPC can implement them using existing vendors or preferred platforms.

7.1 Identity Hardening (Highest Priority)

- Enforce MFA for privileged and remote access across AD, cloud consoles, and CI/CD [10], [8], [9].
- Implement PAM (just-in-time elevation) and separate admin accounts/workstations.
- Apply least privilege and remove broad local admin rights; monitor privileged logons centrally [3].

7.2 Network and Service Exposure Reduction

- Reduce exposed services and enforce allow-lists using host firewalls and security groups [8].
- Segment portal, CI/CD, and R&D identity tiers; restrict management protocols to admin networks.
- Add egress controls and DNS policy to limit covert command-and-control and exfil paths.

7.3 Application and CI/CD Security Controls

- Adopt a secure SDLC: SAST/DAST, dependency scanning, and secret scanning in pipelines [9].
- Enforce code review, branch protections, and least-privilege CI tokens; rotate secrets routinely.
- Require artefact provenance/signing for release builds and monitor anomalous pipeline executions.

7.4 Centralised Detection and Response

- Centralise identity, endpoint, and cloud telemetry in a SIEM and align detections to ATT&CK [1].
- Create playbooks for scanning, suspicious authentication, and web exploitation indicators [4].
- Measure time-to-detect and time-to-contain as key performance indicators.

7.5 Security Validation as a Continuous Practice

- Run periodic purple-team exercises and regression tests after significant changes [2].
- Maintain an external attack surface inventory and patch governance cadence.
- Track remediation via a risk register with clear owners, deadlines, and verification.

8. Cost-Benefit Analysis (Business Case)

For pharma companies, the risks of IP theft are significant and can include loss of competitive advantage, regulatory implications, reputational harm and delayed product release. As such, CPC should focus on the controls that decrease the probability of compromise and time to detect. Example cost-benefit comparison Table 3 illustrates a hypothetical benefit-cost comparison; these values are indicative only and should be used for the purpose of prioritising rather than purchasing. For CPC, they become more valuable when changes are associated with measurable

KPIs (e.g., dwell time, reduction of privileged access and attack surface on the portal) and reviewed by governance on a quarterly basis.

Investment	Indicative Cost (Year 1)	Primary Benefit	Why it Matters for IP Theft
MFA + Conditional Access + PAM	Medium	Reduces credential abuse and privilege escalation	Identity is the control plane for AD, cloud, and CI/CD.
EDR + Central SIEM Integration	Medium-High	Improves detection and containment speed	Shortens dwell time and limits collection/exfil.
Segmentation + Firewall/SG Hardening	Medium	Constrains lateral movement and blast radius	Prevents portal compromise from reaching R&D/AD.
Secure SDLC in CI/CD (SAST/DAST/Secrets)	Low-Medium	Prevents shared web and supply-chain flaws	Reduces initial access vectors through the portal/pipeline.
WAF + App Monitoring	Low-Medium	Blocks and detects malicious web patterns	Detects injection and parameter tampering attempts early.

9. Conclusion

Operation IP Shield notes that this hybrid exposure of CPC is best reduced by robust identity controls, minimisation of exposed services, and end-to-end telemetry. The data demonstrates that recon and app-level fuckery can rapidly escalate if your patching/input validation/monitoring is uneven. By adopting MFA/PAM, segmentation, secure CI/CD, and ATT&CK-aligned detections via a centralised SIEM/EDR stack, CPC will significantly reduce the likelihood of IP loss and accelerate time-to-detect and time-to-contain.

Appendix A. Evidence

The following figures reproduce the captured evidence screenshots.




Figure 1. ICMP reachability verification between the Kali workstation and the target host.




Figure 2. A baseline network scan was executed to enumerate open ports and exposed services on the target.




Figure 5. Aggressive scan (-A) capturing service detection and script output to enrich reconnaissance.




Figure 6. Aggressive scan (-A) continuation showing additional enumerated services and OS hints.




Figure 7. Aggressive scan (-A) continuation showing further service enumeration output.




Figure 8. Aggressive scan (-A) continuation showing additional reconnaissance details.




Figure 9. Aggressive scan (-A) finalisation showing host/service inventory after deep probing.




Figure 10. Metasploit module search results for a service-specific exploit candidate (ProFTPD) during exploit selection.




Figure 11. Exploit workflow attempt in Metasploit demonstrating configuration of remote host parameters and execution trial.




Figure 12. Follow-on exploit attempt illustrating error handling and validation of required module options (target specification).




Figure 13. Metasploit evasion options view is used to review available protocol-level obfuscation settings




Figure 14. Payload compatibility list used to select an appropriate session payload for controlled testing.




Figure 15. Successful lab compromise indicator and session establishment evidence for the authorised target.




Figure 16. Post-compromise validation output confirming shell access and host context in the lab environment.




Figure 21. DVWA authentication page used to validate basic access control behaviour with a test account (baseline).




Figure 22. Burp Suite intercepted HTTP request packet for DVWA login, captured for inspection and analysis.




Figure 23. Burp Suite inspector view used to modify request body parameters for controlled testing




Figure 24. Burp Suite modification state illustrating an injection-style test payload in the request (lab validation).




Figure 25. Resulting modified request after parameter adjustment, used to observe server-side validation and behaviour.




Figure 26. Final stage of the request manipulation sequence showing the crafted request sent for controlled test execution.




Figure 27. Reachability test (ping) from Kali to the on-premises Active Directory host (hybrid environment connectivity check).




Figure 28. Metasploit database/service workflow initialisation (msfdb/ workspace start) prior to importing scan results.




Figure 29. Nmap [12] scan output highlighting discovered Windows/AD-related services on the domain host (AD surface mapping).




Figure 30. Service/version detection output used to infer OS family and key services




Figure 31. Aggressive scanning output aimed at extracting domain context




Figure 32. Aggressive scanning continuation showing host script results




Figure 33. DNS-oriented enumeration test results used to validate name service exposure and potential misconfiguration signals.

References

- [1] MITRE, “MITRE ATT&CK®: Enterprise Matrix,” 2025. [Online]. Available: <https://attack.mitre.org/matrices/enterprise/>. Accessed: Dec. 17, 2025.
- [2] National Institute of Standards and Technology (NIST), “Technical Guide to Information Security Testing and Assessment,” NIST Special Publication 800-115, Sep. 2008. [Online]. Available: <https://csrc.nist.gov/publications/detail/sp/800-115/final>. Accessed: Dec. 17, 2025.
- [3] NIST, “Security and Privacy Controls for Information Systems and Organizations,” NIST Special Publication 800-53 Rev. 5, Sep. 2020. [Online]. Available: <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>. Accessed: Dec. 17, 2025.
- [4] OWASP Foundation, “OWASP Top 10 2021: The Ten Most Critical Web Application Security Risks,” 2021. [Online]. Available: <https://owasp.org/Top10/>. Accessed: Dec. 17, 2025.
- [5] DVWA Project, “Damn Vulnerable Web Application (DVWA),” GitHub repository. [Online]. Available: <https://github.com/digininja/DVWA>. Accessed: Dec. 17, 2025.
- [6] PortSwigger, “Burp Suite Documentation: Proxy,” 2025. [Online]. Available: <https://portswigger.net/burp/documentation/desktop/tools/proxy>. Accessed: Dec. 17, 2025.
- [7] Rapid7, “Metasploit Framework Documentation,” 2025. [Online]. Available: <https://docs.rapid7.com/metasploit/>. Accessed: Dec. 17, 2025.
- [8] Amazon Web Services, “Amazon VPC Documentation,” 2025. [Online]. Available: <https://docs.aws.amazon.com/vpc/>. Accessed: Dec. 17, 2025.
- [9] GitHub, “Security hardening for GitHub Actions,” 2025. [Online]. Available: <https://docs.github.com/en/actions/security-guides/security-hardening-for-github-actions>. Accessed: Dec. 17, 2025.
- [10] Microsoft, “Active Directory Domain Services overview,” 2025. [Online]. Available: <https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/get-started/virtual-dc/active-directory-domain-services-overview>. Accessed: Dec. 17, 2025.
- [11] E. M. Hutchins, M. J. Cloppert, and R. M. Amin, “Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains,” Lockheed Martin Corp., white paper, 2011.
- [12] G. F. Lyon, Nmap Network Scanning: The Official Nmap Project Guide. Insecure.Org LLC, 2009. [Online]. Available: <https://nmap.org/book/>. Accessed: Dec. 17, 2025.