



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Ethical Hacking

ACCA7033 Accessment Report

Student Name: Gandhi Raj Rakeshkumar

Student ID: 2503898

Program: MSC Cybersecurity and Digital Forensics

Submission To: Dr.Mohamed Ghanem

Table of Contents

1. Introduction.....	1
2. Background of the Incident.....	1
3. ATT&CK Mapping and Analysis MITRE ATT.....	2
3.1 Reconnaissance	2
3.2 Initial Access	2
3.3 Execution	2
3.4 Persistence	2
3.5 Credential Access and Privilege Escalation	2
3.6 Lateral Movement	3
3.7 Command and Control	3
3.8 Collection and Exfiltration	3
4. Root Cause Analysis	3
4.1 Security Control Failures	3
4.2 Architectural Weaknesses	3
4.3 Human and Organisational Factors	4
4.4 Security Shortcomings in the Clouds	4
4.5 Summary	4
5. Strategic Recommendations.....	4
5.1 Preventive Security Controls	4
5.2 Architectural Enhancements	5
5.3 Detection and Monitoring Enhancements	5
5.4 Strengthening Cloud Security	5
5.5 Organisational and Human Controls	5
6. Literature synthesis, Academic Reflection	6
7. Strategic Insight Wrap-Up	6
8. Introduction to Red Team Engagement	6
9. Rules of Engagement (RoE)	7
9.1 Scope Definition	7
9.2 Operational Boundaries	7

9.3 Authorised Techniques	7
9.4 Reporting Confidentiality Protocols	7
9.5 Legal and Ethical Factors	7
10. Threat Emulation Plan	8
10.1 Objective Alignment	8
10.2 Attack Phases and MITRE Mapping	8
10.3 Operational Security (OPSEC) Considerations	9
10.4 Success Criteria	9
11. Environment Design and Architecture	9
11.1 Infrastructure Overview	9
11.2 On-Premises Environment	9
11.3 Cloud Infrastructure	9
11.4 CI/CD and Developer Stack	10
11.5 Network & Interconnectivity Diagram	10
12. Attack Execution Narrative	10
13. Detection Gap Analysis	21
13.1 Blind Spots to Identity and Access	21
13.2 Weaknesses of Endpoint Detection	21
13.3 Network and Lateral Movement Visibility Problems	21
13.4 C2 and Exfiltration Monitoring Shortcomings	22
13.5 Summary	22
14. Strategic Recommendations	22
14.1 Security Improvements of Identity	22
14.2 Architecture and Improvement Segmentation	22
14.3 Better Detection and Surveillance	22
14.4 Cloud Scale and Hardening of CI/CD	22
14.5 Human and Operational Controls	22
15. Cost–Benefit Justification	23
16. Executive-Level Summary	23
17. Conclusion	23
References	24

1. Introduction

The rising rates of occurrence and level of sophistication of a series of cyber intrusions into pharmaceutical and healthcare organisations have given rise to a surge in the concentration on the protection of intellectual property and resilience to cyber-attacks at the national level. One of the most well-planned efforts to compromise the COVID-19-related research settings took place during the global pandemic, and it is the campaign associated with APT29, or Cosy BeCosy. The attacker, who is thought to have acted under the guidance of a state, is an espionage-oriented attack against the information regarding the development of vaccines, proprietary biomedical modelling, and confidential clinical trial data. The operation was expert in terms of operations that showed a high degree of maturity, stealth, zero-day attacks, advanced malware delivery, and adaptive command and control techniques. The incident is a valuable case study in exploring the aspects of adversarial capability, detection gaps, and resiliency failures within the hybrid research setting. Interpreting this operation within the framework of a well-organised threat intelligence process allows identifying it with the established models, including the MITRE ATT&CK Enterprise Matrix, and facilitates the strategic planning of defence.

2. Background of the Incident

In July 2020, the National Cyber Security Centre (NCSC) of the United Kingdom, the Cybersecurity and Infrastructure Security Agency (CISA) of the United States and the Canadian Communications Security Establishment (CSE) officially attributed APT29 to a coordinated cyber espionage program. The organisations of interest were pharmaceutical companies, labs that produced vaccines and academic medical research institutions that worked on the development of SARS-CoV-2 vaccines and therapeutics. Official notices assured that the main target would be obtaining valuable research and not interference, financial blackmail, or sabotage.

APT29 used a multi-phase intrusion that included enacting specific reconnaissance, which was then seen through acquiring credentials and exploiting publicly accessible infrastructure. The initial access points were spear-phishing mailings with personalised messages and victimised infrastructure [1]. The team used malware families, which included WellMess and WellMail, which were specifically created to run with minimal forensic footprint and integrate into normal processes. Leverage of vulnerabilities, such as CVE-2019-19781 (Citrix ADC) and CVE-2020-0688 (Microsoft Exchange vulnerability), was an increased opportunity of access, with phishing alone not being sufficient to do so.

With the network accessible, persistence methods involved altered registry keys, tasks, and encrypted data communications of the HTTPS protocol and at times cloud-based infrastructure. Exfiltration was aimed at the controlled staging of data and, reduction of outbound abnormalities. The operation went unnoticed in various settings over time, clearly demonstrating both advanced tradecraft and defensive control failure [2].

3. ATT&CK Mapping and Analysis MITRE ATT

3.1 Reconnaissance

APT29 opened operations with specific information collection based on a MITRE ATT&CK methodology of T1592 (Gather Victim Host Information) and T1593 (Search Open Websites/Domains). Intelligence footprint included publications in peer-reviewed scientific journals, lists of speakers at conferences, LinkedIn accounts and profiles, as well as patterns of e-mail among employees. The group also used T1598.002 (Phishing for Information: Spear-phishing Service) in order to authenticate credentials ahead of an active intrusion event.

3.2 Initial Access

The first attempt at initial entry was based on T1566.001 (Spear-phishing Attachment), and relied upon employing COVID-19-related baits crafted in such a way that they make them seem convincing to vaccine laboratory staff [3]. Fallback access was used to notice vulnerabilities to T1190 (Exploit Public-Facing Application) when the email delivery controls were functioning as intended. Use of Microsoft Exchange and Citrix ADC vulnerabilities offered an authenticated access and multi-factor authentication bypass where misconfigurations occurred.

3.3 Execution

The implementation was based on the execution of evil payloads with T1059.001 (PowerShell) and T1204 (User Execution) mapped. Malicious DLLs were executed by WellMess (custom malware) and encrypted download-and-execute tasks were executed as well (instead of T1220.001) using T1218.011 (Signed Binary Proxy Execution: Rundll32) to provide stealth and anti-forensic behaviour.

3.4 Persistence

The mechanisms of persistence were consistent with T1053.005 (Scheduled Task/Job) and T1547.001 (Registry Run Keys). Further persistence of clouds came in the form of compromised Microsoft 365 service accounts under T1136.003 (Create Cloud Account), in which the identity federation policies were not enforced.

3.5 Credential Access and Privilege Escalation

Privilege escalation occurred as a result of a mixture of T1055 (Process Injection) and T1068 (Exploitation for Privilege Escalation). Credential harvesting was through various protocols, T1003.006 (LSASS Memory Dumping) and by using OAuth tokens in the cloud platforms. The internal movement was highly enhanced in settings where there was no PAM governance in place [4].

3.6 Lateral Movement

Such inter-system movement was in accordance with T1021.001 (Remote Services: RDP) and T1071.004 (DNS) to have covert C2 traffic. The methods of Kerberos manipulation and pass-the-hash were in line with T1550.002 (Use of Stolen Credentials), which makes it possible to traverse segmented research and administrative backgrounds.

3.7 Command and Control

C2 infrastructure also adhered to the T1071.001 (Web Protocols) based on HTTPS masquerading as bgn browser traffic. OPEC countermeasures, Dynamic domain name and TLS certificate rotation supported OPSEC as per T1008 (Fallback Channels).

3.8 Collection and Exfiltration

The staging of targeted data took place under the terms of T1074.001 (Local Data Staging), and the ultimate exfiltration followed the terms of T1048.002 (Exfiltration over Asymmetric Encrypted Channels). Cloud drive synchronisation abuse was a secondary channel that conformed with T1567.002 (Exfiltration to Cloud Storage).

Summary

The operation showed high confidence of operational maturity, adversarial stealth operational capabilities, and alignment with the strategic objectives of nation state accumulation of intelligence.

4. Root Cause Analysis

With an in-depth analysis of the APT29 case, it is evident that it had numerous contributory factors at all levels in the technological, architectural, and human levels. The breach was caused by neither a single point of failure, but a stratified crumpling of the preventive, detective and response methods [5]. Several systemic flaws provided favourable environments of adversarial penetration, privilege escalation and controlled exfiltration.

4.1 Security Control Failures

Insufficiency of enforced multi-factor authentication on all remote services facilitated the use of credentials and access after phishing and harvesting of the credentials. This coincides with identity governance controls failure where hybrid Microsoft Exchange and Microsoft 365 are present. Moreover, endpoint detection tools did not provide behavioural baselining of command execution, especially PowerShell abuse, to support propagation, execution and credential looting operations without creating significant alerts [6]. A gap in logging also contributed to the absence of early detection, showing inadequate monitoring of threat behaviours associated with MITRE ATT&CK.

4.2 Architectural Weaknesses

The hybrid infrastructure had a lack of segmentation between research issues, the administrative domain and services available to the outside world. After obtaining authenticated access, the way lateral movement paths existed was as a result of a flat twork trust relationship. Exposure was also

supported by legacy systems with old authentication models, especially when there is a limitation of imposition on conditional access controls due to compatibility obligations. Depending on the services that are open to the external world, like C, Irix ADC and Microsoft Exchange have also broadened the attack path, thereby increasing the likelihood of exploiting the services by known or novel vulnerabilities.

4.3 Human and Organisational Factors

The carelessness of humans, insufficient long-term time awareness regarding cyber exposure, provided a vulnerability as a social engineering vector [7]. The developers of vaccines were focusing more on speed of operations rather than adherence to protocols and particularly in the conditions of a pandemic, when remote access grew exponentially. This setting generated less scrutiny of online communications and negligence towards high-phishing click-through probability. Also, the preparedness to act in response to incidents was not sufficient; the adversary ensured continuity over long periods of time since it took too long to detect the anomaly and covered the detection scope and complete the process of escalation.

4.4 Security Shortcomings in the Clouds

Weak policy setups in identity federation, lack of cohesive cloud access control enabled attackers to hold to ground using your stolen credentials in Microsoft 365. The token persistence was so that an event of password reset could allow long-term access indicating that the token revocation and session lifecycle management are weak [8].

4.5 Summary

The root cause analysis finds that the breach to take place was caused by a combination of a lack of identity security, inadequate network segmentation, open services to the Internet, limited behavioural planning, and human cyber hygiene failures. The opponent took advantage of this combination to have a sustained, discrete access and attain its intelligence-gathering goals.

5. Strategic Recommendations

5.1 Preventive Security Controls

This should be strengthened by focusing on identity security by using phishing-resistant multi-factor authentication like FIDO2 or certificate-based authentication. Privileged Access Management (PAM) will have to be implemented to control administrative user identities, implicitly enforce credential vaulting, and prohibit credential harvesting. There should be application allow-listing to stop unacceptable running of LOLBins like PowerShell, mshta, and run, dll32 beyond sanctioned behavioural patterns [9].

The next phase in endpoint security should be Extended Detection and Response (XDR) with behavioural analytics, anomaly detection, and MITRE ATT&CK mapping of rules. Vulnerability management needs to continuously focus on the in-depth nature of high-severity exploitation vulnerabilities that public-facing organisations have.

5.2 Architectural Enhancements

Zero Trust Architecture (ZTA) is a redesign strategy requirement. The network should be segmented to identify clinical research systems, authentication structure, and cloud services. This should be granted with least privilege, adaptive trust scoring and just-in-time authorisation. Lateral movement can be blocked later through the use of software-defined security controls that implement micro-segmentation. The policies of conditional access must implement device compliance, geographical location limitations and identity risk scoring.

The old authentication schemes that include NTLM should be phased out and only NTLM and certificate-based should be used [10]. The authorisation configuration baselines and automated compliance auditing across Windows, Linux and cloud infrastructure workloads should be provided.

5.3 Detection and Monitoring Enhancements

SIEM and XDR systems also need to be co with aggregated endpoint telemetry, identity provider telemetry, as well as Active Directory telemetry, cloud telemetry, and CI/CD pipelines. The detection engineering based on MITRE must generate detection signatures in line with important adversarial behaviours that cover credential dumping, registry persistence, scheduled task creation, cloud token replay and encrypted beaconing.

Proactive testing model AtT&C that should be formalised to threat hunting programmes should be ATT&C. This should be countered by DNS anomaly surveillance, privacy checks on encrypted traffic, and the identification of beacon frequency to combat the covert C2 communication.

5.4 Strengthening Cloud Security

Cloud services access should be conditional and feature continuous authentication, token lifetime restrictions, device compliance being mandatory and automatically revoking sessions. Cloud workload protection platforms (CWPP) and cloud security posture management (CSPM) must be integrated to find out misconfigurations and unauthorised access to the storage, identity, and cloud resources.

5.5 Organisational and Human Controls

The culture of security should be concentrated by conducting unceasing phishing simulation training, obligatory secure research handling behaviorsbehaviourster-of-role cyber training to clinical researchers, laboratory staff and administrative users. An incident response playbook that is crisis-ready and in accordance with the NIST and ISO paradigms needs to be developed and exercised on a regular basis.

Summary

These recommendations are a multi-layered defensive plan that deals with identity security, architecture, monitoring, and governance and human-friendly resilience. The issue of focus is the shift towards active deployment of control to preemptive adversary understanding of defence.

6. Literature synthesis, Academic Reflection

Recent academic and industrial studies indicate that there is a definite intersection point between motivating factors in the state-aligned threat actors and a growing strategic importance of biomedical and pharmaceutical data [11]. The incident with APT29 essentially compares with the larger academic results that healthcare and research fields are highly desirable alone for espionage goals because of the intellectual property concentration and historically lower cyber maturity than other economic or defence fields.

An APT tradecraft literature analysis reveals that there is a strict adherence to the use of the crthe edentialed access, cloud persistence and living-off-the-land methods to minimise forensic signature. These trends are confirmed in news on MITRE Engenuity, CrowdStrike Global Threat Intelligence and ENISA threat surveys, which all indicate that the present-day defences need to include behaviour-based surveillance and zero-trust centred on identity models and not on a sole perimeter or signature-driven security.

Research exploring organisational readiness to respond towards technological threats suggests that technological controls can only be effective with tangible Incident response governance maturity, cultural reinforcement and documented Incident response systems [12]. Literature also shows that hybrid and cloud environments necessitate innovative types of visibility and identity lifecycle controls to minimise access threats, as seen in the case.

7. Strategic Insight Wrap-Up

The discussion of the APT29 intrusion demonstrates the extreme significance of integrating the security operations with adversary-informed intelligence. This event has proved that the conventional perimeter-driven defences cannot resist the stealth-related and credential-oriented attacks. The identified mapped behaviours, root causes, and strategic suggestions are all pieces of a coherent structure of enhancing resilience in the hybrid environments. Constant monitoring, Zero Trust, enhanced identity controls, and MITRE-compatible threat hunting become the key foundations of the current cyber defence [13]. The event justifies the proactive and intelligence-focused security maturity in comparison to the reactive patching or deployed isolated controls in place to guarantee being ready to respond to the next nation-state level attacks in the future.

8. Introduction to Red Team Engagement

Operation IP Shield is an active red-team action that will test the resilience of the hybrid infrastructure of Cureton Pharmaceutical Company to an emulated Advanced Persistent Threat. The campaign goal is the successful full acquisition of secured research data in an undocumented way, which reflects realistic threat behaviour consistent with the state. The scope of engagement comprises an on-premises Active Directory environment, an AWS-hosted data platform, and built-in pipelines employed in drug development systems. The operation will be organised of the MITRE ATT&CK framework to facilitate all the threat-focused execution, reporting, and defensive learning outcomes. The results will lead to strategic risks treatment and improved cyber defence maturity.

9. Rules of Engagement (RoE)

9.1 Scope Definition

The approved scope of the Operation IP Shield comprises the internal Active Directory domain of Cureton Pharmaceutical Company alongside the cloud computing AWS VPC that contains and processes clinical research data, and the CI/CD processes combined with GitHub Actions [14]. Everything in these environments qualifies as part of the in-scope items to be reconnoitred, escalated privileges, testing persistence, and extraction of controlled data. Some of the out-of-scope assets are the production medical devices, the personal devices of employees, and the financial systems that are not linked with the research activities.

9.2 Operational Boundaries

Activities in testing assume non-destructive and reversible testing. Denial-of-service attack, intentional corruption of system settings and the deployment of destructive payloads are forbidden. Only accepted payload types, implants and methods of command execution discovered in the engagement briefing can be. The payloads can not exceed accepted execution rights and may not alter valuable operational parts of any system other than that which is necessary to validate the persistence or the theming of sideways movements.

9.3 Authorised Techniques

Engagement permits the use of phishing emulations, harbour credentials, existing vulnerabilities or hacks, use of living-off-the-land binaries (LOLbins), encrypted command-and-control networks and forms of identity abuse on clouds. The ways of physical access, hardware manipulation, wireless intercept, and social engineering are not allowed towards external partners.

9.4 Reporting Confidentiality Protocols

There is a managed communication system that controls the execution of operations. A nominated oversight officer has the full discretion to stop or end activities based on the stipulated termination phrase: "Suspend Vector. Progress checkpoints should be recorded in the red-team operational ledger weekly, and there should be a record of the actions, systems accessed, persistence deployed and events of data handling. Any system event or undesired stability causes automatic suspension of the active tooling and filing of an anomaly report.

9.5 Legal and Ethical Factors

Every activity is under the law of national cybersecurity, organisational compliance policies and professional ethics requirements. The involvement is done with the involvement of the executive and is restricted to authorised personnel [16]. Only approved dummy research data can be extracted out of fromta and stored in the encrypted evidence repository stated in the engagement directive.

10. Threat Emulation Plan

10.1 Objective Alignment

The emulation campaign aims at replicating realistic APT behaviour against pharmaceutical intellectual property. Scenarios feature activities to quantify the detection capability, architectural resilience, identity governance enforcement, and readiness to respond.

10.2 Attack Phases and MITRE Mapping

The plan aligns operational actions with the following phases:

Phase 1: Reconnaissance

Activities include harvesting public information aligned with T1593 (Search Open Websites/Domains) and passive DNS analysis. Identification of exposed cloud assets, employee credential formats, and research-related digital footprints forms the targeting baseline.

Phase 2: Initial Access

Two primary entry vectors will be attempted:

- Spear-phishing (T1566.001) delivering a controlled macro payload
- Public-facing vulnerability exploitation (T1190) targeting AD FS or VPN services

Payload delivery is configured to trigger a remote command channel without overt behavioural indicators.

Phase 3: Execution

Execution relies on:

- T1059.001 PowerShell
- T1218.011 Rundll32 for proxy execution
- T1204 User execution dependency, where required

Execution behaviour is masked using signed binary proxy techniques and in-memory execution tooling.

Phase 4: Persistence and Privilege Escalation

Persistence mechanisms include:

- Scheduled Tasks (T1053.005)
- Registry Run Keys (T1547.001)
- Cloud token persistence (T1098.003)

Privilege escalation will utilise credential dumping (T1003.001) and exploitation of misconfigured privileges (T1068).

Phase 5: Lateral Movement

Movement between systems leverages:

- Pass-the-hash (T1550.002)
- Remote Desktop Protocol (T1021.001)
- SMB session reuse (T1021.002)

Target priority includes domain controllers, research data staging servers, and CI/CD orchestrators.

Phase 6: Data Collection and Exfiltration

Data staging aligns with T1074.001 (Local Staging). Exfiltration mechanisms include:

- HTTPS encrypted channels (T1041)
- Cloud exfiltration via existing accounts (T1567.002)

Only dummy datasets will be extracted, and the extraction size must remain below anomaly thresholds.

10.3 Operational Security (OPSEC) Considerations

Noise reduction measures include beacon jitter randomisation, sleep cycles, and avoidance of high-risk commands. Logging evasion aligns with T1562.002 Disable Security Tools where permitted.

10.4 Success Criteria

Success is defined by:

- Covert C2 channel establishment
- Verified persistence
- Undetected lateral movement
- Successful exfiltration of the approved dataset

11. Environment Design and Architecture

11.1 Infrastructure Overview

The hybrid laboratory that was applied during operation IP Shield is a duplicate of the production system of Cureton Pharmaceutical Company [18]. The environment comprises three parts that are both integrated: an on-premises Active Directory(AD) domain, used to authenticate different users within the company, is based in a Virtual Private Cloud (VPC) within the AWS data infrastructure, which holds clinical research data, and finally CI/CD infrastructure is used to support development processes in the form of GitHub Actions.

11.2 On-Premises Environment

The physical environment consists of an on-premises environment with a domain controller based on Windows Server, user workstations, an internal file server and a gateway accessible via VPN. Intellectual grouping Logical grouping involves research endpoints that are assigned a specific organisational unit, which is indicative of sensitive intellectual property held within the organisation [19]. It has limited network segmentation, allowing authorised traffic between infrastructure nodes and the user subnets. Access to administration is by use of AD group policies.

11.3 Cloud Infrastructure

The clinical data platform provided by AWS resides in a segmented VPC that has a public subnet and a private subnet. The facade layer will have an application load balancer and web application firewall that support secure access to the total data portal. The private subnet comprises the S3 storage repositories, EC2 compute instances to process the data, and role-based controls in IAM. Microsoft 365 and AWS are identity federated to provide single sign-on, which represents an enterprise hybrid identity model.

11.4 CI/CD and Developer Stack

This pipeline development incorporates the GitHub Actions and the containerised workloads. Secrets administration and environment variables are kept in the encrypted vault of GitHub, which resembles the real-world DevOps pipelines [20]. Build artefacts employed in the AWS environment for forming connections between development and production ecosystems.

11.5 Network & Interconnectivity Diagram

The architecture comprises VPN tunnels, plus identity federation trust among all three environments. This provides practical cloud as well as on-premise inter-availability of lateral movement. Logging is somewhat centralised into an instance of SIEM, but there are not many correlation rules.

12. Attack Execution Narrative

Phase 1: Reconnaissance

The initial operation was passive reconnaissance of publicly available data concerning the infrastructure of Cureton Pharmaceutical Company. Open-source intelligence tools found employee email format, revealed cloud services, and a system of clinical data management on T1593 (Search Open Websites/Domains). Other intelligence used on metadata, job advertisements, GitHub accounts, and LinkedIn accounts showed the presence of Microsoft Active Directory, VPN remote access services, Microsoft 365 identity integration, and GitHub Actions-based CI/CD pipelines [21]. This step formed an awareness of the internal technologies, naming standards, and possible attack points without causing any type of defensive alarms.

Phase 2: Initial Access

The main entry point of the breach was a directed spear-phishing email and was consistent with T1566.001 (Spear-phishing Attachment). A rogue macro-enabled document, which appeared as a research compliance form, was a payload loader scaled on a Webservrweb server that would execute an encrypted command as it is opened. After being triggered, the loader found a connection to the specified command-and-control (C2) infrastructure that used HTTPS and added a Sliver beacon. Simultaneously, a different preliminary access approach consulted a vulnerable VPN endpoint, and a recognised configuration matched the T1190 (Exploit Public-Facing Application). The secondary vector was redundancy and not a precondition as obligatory.

Phase 3: Execution

This was executed by in-memory payload deployment as T1059.001 (PowerShell) and T1218.011 (Signed Binary Proxy Execution Rundll32). These tricks are the implementation of shellcode without the need to save executable code. The beacon payload opened up an encrypted persistent channel of communication, which posed as common web traffic on port 443 [22]. Patterns of command execution were deliberately low frequency, borrowing sleep-jitter cycles to prevent behavioural analytics phenomena with respect to bursts of rapid command execution. Stealth was also used by taking advantage of legitimate windows including svchost.exe and explorer.exe, as part of behavioural camouflage.

Phase 4: Persistence

Persistence mechanisms were deployed once access to execution had been confirmed, so as to achieve a long-term, continuous foothold. T1053.005 (Scheduled Task/Job) was used to create a scheduled task by using system-native utilities. An update of T1547.001 (Registry Run

Keys/Startup Folder) to persist as a registry control needed to be executed every time the system began. In the cloud environment session, tokens were compromised, having persistence in line with T1098.003 (Additional Cloud Credentials), which didn't require a password change to proceed with the access. All of the persistence techniques were authenticated by controlled reboots of the system to ensure that it remained operated. Operational.

Phase 5: Privilege Escalation

Privilege escalation applied a variety of techniques depending on the target system. T1068 (Exploitation for Privilege Escalation) was used to determine system hosts that had misconfigured privileges (on-premises Windows). Harvesting of credentials based on T1003.001 (LSASS Memory) gave us user tokens such as research, infrastructure administrators, and DevOps [23]. The harvested administrator tokens were the level of domain escalation as well as the access to essential file shares, group policies, and CI/CD infrastructure. On AWS, an improperly set identity policy allowed privileged access controls that facilitated privilege escalation through IAM enumeration that matched the T1484.002 (Domain Trust Modification)-equivalent cloud tactics.

Phase 6: Lateral Movement

The lateral movement had a strategic path that focused on exposure minimisation. With the help of harvested credentials, access was redirected to the research subnet with T1021.001 (Remote Desktop Protocol) and T1550.002 (Use of Stolen Credentials). Secondary access channels were PowerShell and WinRM. The lateral movement mapping was based on visibility of network shares, queries to domain controllers and stores that contained CI/CD credentials. Navigation did not attempt any internal scanning bursts in order to lower the chances of a defensive tooling anomaly prompt being detected.

After the identity re-use was confirmed, attacks were allowed on AWS Cloud workloads using compromised Microsoft 365 federated identities. OAuth tokens were used to gain access to CI/CD repositories and container-based workloads used in the automation of pipelines. The pivoting was entirely credential-based and did not make use of exploitation tools that can instigate a high-severity alert [24].

Phase 7: Data Collection


These collection activities were in harmony with T1074.001 (Local Data Staging). Dummy intellectual property assets such as encrypted clinical research logs, proprietary analysis models stored on file shares and cloud storage buckets were found. Part of staging directories in legitimate working paths like 0,x DATAP: and cloud-backed folders were used in order to reduce the element of detection risk. The approved dataset was only gathered and temporarily compacted using native windows. Windows that file signature anomalies would not exist.

Phase 8: Command and Control



Communications were sustained throughout the entire phases through an encrypted Silver T1071.001 (Web Protocols). Plausible traffic patterns were guaranteed by beacon jitter, variable sleep intervals and domain fronting [25]. DNS fallback channels in line with T1071.004 (DNS) were provided as alternate channels. C2 infrastructure turned over TLS certificates on a regular basis to reflect commercial domains.

Phase 9: Exfiltration



The staged dataset was staged out via encrypted transfer of data via HTTPS, as predicted by T1041 (Exfiltration Over Command-and-Control Channel). To assume normal browsing traffic, it slowed down its transfer rate, which prevented sharpness that is alarming to traffic monitoring devices. Cloud exfiltration based on synchrosynchronised under T1567.002 (Exfiltration to Cloud storage) was validated and not activated on redundancy. After successful exfiltration was achieved,







[illegible][illegible]




[illegible][illegible]









13. Detection Gap Analysis

13.1 Blind Spots to Identity and Access

Lack of phishing-resistant multi-factor authentication and lax access enforcement allowed reusing credentials without a warning notice. The misuse of cloud session tokens was not detected since it lacked lifecycle governance and behavioural identity analytics.

13.2 Weaknesses of Endpoint Detection

The use of in-memory execution, PowerShell abuse or signed binary proxy methods was not adequately enabled in endpoint logging. Credential dumping and Silver execution had worked without raising defensive responses due to the absence of behaviour-based detection.

13.3 Network and Lateral Movement Visibility Problems

Weaknesses in network segmentation and unlimited access to domains and credentialed relationships with a trust helped the traversal based on credentials without alert information. The logs that were related to the successful RDP connections and remote executions were not correlated with others, which hindered the detection of lateral movement [25].

13.4 C2 and Exfiltration Monitoring Shortcomings

Outbound C2 encrypted over HTTPS was very similar to legitimate outbound communication. It did not have TLS scanning of the environment, DNS anomaly scanning and beacon-pattern scanning. The issue of exfiltration throttling and data staging was not detected because there was little integration of data loss prevention (DLP) and a lack of traffic watermarking.

13.5 Summary

The main causes of the detection failures were the lack of identity governance, behavioural telemetry, incomplete integration of SIEM-XDR, and a missing identity telemetry.

14. Strategic Recommendations

14.1 Security Improvements of Identity

The improvement of security posture should start with the enforcement of phishing-resistant authentication using certificate-based security keys. The policy of conditional access must include compliance checking of devices, detection of impossible travel, and permanent authentication.

14.2 Architecture and Improvement Segmentation

Hybrid systems should also take into consideration zero principles. Domain controllers, CI/CD workloads and research databases must be isolated through network micro-segmentation. Cloud access gateways ought to take the shape of private and not public gateways. Credential replay should be eliminated by the removal of legacy authentication protocols.

14.3 Better Detection and Surveillance

The implementation of a detection engineering programme based on the MITRE ATT&CK framework should be proposed [29]. SIEM-XDR association should contain behavioural analytics, anomaly-based and cloud telemetry ingestion. TLS inspection, DNS monitoring and beacon periodicity analysis are required to identify covert C2 channels.

14.4 Cloud Scale and Hardening of CI/CD

CSPM and Cloud Workload Protection Platforms (CWPP) are expected to be implemented to achieve policy compliance and identify misconfigurations. CI/CD pipelines need the sec rotation, zero-trust service modernisation artifact integrity validation [26].

14.5 Human and Operational Controls

The areas that a sustained security training programme should discuss are focused phishing, incident reporting behaviour, and secure data handling of research. Tabletop drills and red team / blue team research must be incorporated in order to confirm preparedness.

15. Cost–Benefit Justification

Security investment should be evaluated in comparison with operational risk and the consequences of a breach. A Pharmaceutical R&D environments possess environments that are rich with highly monetisable intellectual property; compromises may result in competitive disadvantage, regulatory sanctions, loss of market valuation and long-term operational disruption. The estimated cost of an APT-level breach is higher than a strategic safeguard implementation cost.

Trust segmentation, identity security, CSPM/XDR tooling, and incident preparedness features need expenditure on initial capital apparatus and decreasesdecreasebreaches and the duration occupied by the breach [27]. The long-term operational payoff incorporates decreased insurance charges, regulatory adherence, increased investor self-esteem and greater resilience.

16. Executive-Level Summary

The analysis indicates that Cureton Pharmaceutical Company is under legitimate threats of engineered attackers against research and intellectual assets. The operation of the red-teamed team to recreate an intrusion in the style of a nation-state, gaining a backdoor without being noticed, creating persistence, and stealing guarded research information without activating detection in place. The results verify the frailty of identity security, separation of the network, cloud administration, and observation potential [28]. The exposure will be high because of the reinforced security by the use of Zero Trust architecture, increased detection engineering, phishing-resistant authentication, and enhanced incident preparedness. The financial benefits of these actions are compliance with regulations, sustaining business operatioperationsaining reputation, and safeguarding against intellectual property, which is important in competitive advantage.

17. Conclusion

The testing proves the ability of the advanced adversaries to perform hybrid environment research using credential-based intrusion methods, silent persistence, and cloud-based exfiltration methodologies. The APT29 campaign under analysis confirms a larger tendency: the current attacks go by the olaroundperimeter protection measures and are based on deficiencies in identity protection and supply-chain vulnerabilities and the misconfiguration of clouds. The emulation with the red teamthat Cureton Pharmaceutical Company is still susceptible to the credential-based access, horizontal movement, cloud retention, as well as unnoticed data exfiltration. Its results indicate considerable loopholes in monitoring, identity governance and architectural segmentation. The suggested controls can provide an effective roadmap to growing up the seity posture with the transition to intelligence-driven, adversaryadversary-awareareas opposed to the perimeter-based trust.

References

- [1] Z. Lin, “Adaptive cyber defense against APT28: A comparative evaluation of reinforcement learning-based policies in an emulated network environment,” 2025.
- [2] D. Selmanaj, Adversary Emulation with MITRE ATT&CK. Sebastopol, CA, USA: O’Reilly Media, 2024.
- [3] D. Routin, S. Thoores, and S. Rossier, Purple Team Strategies: Enhancing Global Security Posture Through Uniting Red and Blue Teams with Adversary Emulation. Birmingham, U.K.: Packt Publishing, 2022.
- [4] A. Makartsov, “Development of advanced malware action simulator,” Ph.D. dissertation, Vilniaus Gedimino technikos universitetas, Vilnius, Lithuania, 2024.
- [5] J. Hoff, “Creating attack graphs for adversary emulation, simulation and purple teaming in industrial control system (ICS) environments,” FernUniversität in Hagen, Hagen, Germany, 2021.
- [6] E. Granstedt, “Evaluating the realism and effectiveness of automated APT emulation in cybersecurity training: A Lore case study,” 2024.
- [7] K. M. Kareem, The Cyber Eye: Inside the Network Wars and Secrets of the Five Eyes Alliance. Morrisville, NC, USA: Lulu Press, 2025.
- [8] R. Bugler et al., “NATO’s response to hybrid warfare during the Ukraine conflict: Navigating an era of the China–Russia partnership,” 2023.
- [9] E. A. Almasri, “Towards generating a practical SUNBURST attack dataset for network attack detection,” M.S. thesis, Princess Sumaya Univ. for Technology, Amman, Jordan, 2023.
- [10] A. B. Ajmal et al., “Toward effective evaluation of cyber defense: Threat-based adversary emulation approach,” IEEE Access, vol. 11, pp. 70443–70458, 2023.
- [11] A. B. Ajmal et al., “Offensive security: Towards proactive threat hunting via adversary emulation,” IEEE Access, vol. 9, pp. 126023–126033, 2021.
- [12] Z. Zhang, “Reinforcement learning-based approaches for enhancing security and resilience in smart control: A survey on attack and defense methods,” arXiv preprint arXiv:2402.15617, 2024.
- [13] S. Banik et al., “Automated adversary-in-the-loop cyber-physical defense planning,” ACM Trans. Cyber-Physical Syst., vol. 7, no. 3, pp. 1–25, 2023.

- [14] S. Shrivastava, B. Chen, and H. Wang, “DQN learning-based defense against smart primary user emulation attacks in cooperative sensing systems,” *IEEE Access*, vol. 9, pp. 163791–163814, 2021.
- [15] T. Cody et al., “Whole campaign emulation with reinforcement learning for cyber test,” *IEEE Instrum. Meas. Mag.*, vol. 26, no. 5, pp. 25–30, 2023.
- [16] E. Cadet et al., “The role of reinforcement learning in adaptive cyber defense mechanisms,” *Int. J. Multidisciplinary Research and Growth Evaluation*, vol. 2, no. 2, pp. 544–559, 2021.
- [17] V. Colliard, A. Peres, and V. Corruble, “Deep reinforcement learning for dynamic resource allocation: Generalizing defensive strategies based on pairwise unit analysis,” in *Proc. IEEE Conf. Artificial Intelligence (CAI)*, May 2025, pp. 711–718.
- [18] S. Sun, *Redefining Adversarial Dynamics: Co-Evolution of Attack and Defense Strategies in AI-Enabled Power Cyber-Physical Systems*, 2025.
- [19] K. A. Haque et al., “Cyber-physical emulation and threat scenario simulation for enhanced microgrid resilience,” *IEEE Access*, 2025.
- [20] T. Li and Q. Zhu, “Symbiotic game and foundation models for cyber deception operations in strategic cyber warfare,” *arXiv preprint arXiv:2403.10570*, 2024.
- [21] D. R. Bakker, “Autonomous emulation of adversary procedures in the (pre-)compromise domain,” M.S. thesis, Univ. of Twente, Enschede, The Netherlands, 2022.
- [22] J. Wiebe, R. A. Mallah, and L. Li, “Learning cyber defence tactics from scratch with multi-agent reinforcement learning,” *arXiv preprint arXiv:2310.05939*, 2023.
- [23] Y. Gao et al., “Reinforcement learning-based anti-jamming strategy for self-defense jammer-aided radar systems,” *IEEE Trans. Aerosp. Electron. Syst.*, 2024.
- [24] S. H. Oh et al., “Applying reinforcement learning for enhanced cybersecurity against adversarial simulation,” *Sensors*, vol. 23, no. 6, p. 3000, 2023.
- [25] M. Kouremetis et al., “Mirage: Cyber deception against autonomous cyber attacks in emulation and simulation,” *Annals of Telecommunications*, vol. 79, no. 11, pp. 803–817, 2024.
- [26] B. G. Tekgul et al., “Real-time adversarial perturbations against deep reinforcement learning policies: Attacks and defenses,” in *Proc. European Symp. Research in Computer Security (ESORICS)*, Sep. 2022, pp. 384–404.
- [27] S. Lin et al., “BattleAgent: Multi-modal dynamic emulation on historical battles to complement historical analysis,” *arXiv preprint arXiv:2404.15532*, 2024.

- [28] E. Iturbe et al., “Reinforcement learning in action: Powering intelligent intrusion responses to advanced cyber threats in realistic scenarios,” *Expert Systems with Applications*, p. 129168, 2025.
- [29] Z. Zhu et al., “Effective defense strategies in network security using improved double dueling deep Q-network,” *Computers & Security*, vol. 136, p. 103578, 2024.
- [30] M. H. Ashik et al., “Phantom Guard: An MDP-based anti-jamming defense strategy for military 5G networks,” in *Proc. IEEE MILCOM*, Oct. 2024, pp. 1100–1105.
- [31] M. F. Saiyed and I. Al-Anbagi, “A game theoretic model for strategic defence selection against DDoS attacks in IoT networks,” *IEEE Trans. Network and Service Management*, 2025.