OFFENSIVE GUARDIAN

CONFIDENTIAL · SPECIMEN

ADVERSARY VALIDATION · FULL-SCOPE ENGAGEMENT

Adversary Validation Report

Northwind Mutual Assurance Group · objective-led red team operation and defensive chain validation, including the June retest record.

A demonstration document. The client is fictional. The deliverable is the real one.

ENGAGEMENT	ENGAGEMENT WINDOW	RETEST WINDOW
OG-2026-0143	2 Mar – 24 Apr 2026 OPERATION 2 MAR - 11 APR	8 – 12 Jun 2026
ISSUED	**DOCUMENT VERSION**	**ENGAGEMENT LEAD**
19 June 2026	1.0 · Final	Jean-François Maes

Evidence, not assumptions.

OFFENSIVE-GUARDIAN.COM

READ THIS FIRST

Northwind Mutual Assurance Group does not exist. Its estate, its people, and every result in this report were built to show what an Offensive Guardian engagement produces. Hostnames use reserved documentation names (RFC 2606) and addresses use reserved documentation ranges (RFC 5737, RFC 1918). No real system was tested to write it, and no real client's compromise appears in these pages.

What is not invented is the deliverable. The method, the two rating axes, the finding template, the evidence discipline, the improvement plan and the retest record are exactly what a client receives, down to the sections that report what we could not verify.

Read it as an instrument rather than a case study. Because the environment was constructed rather than encountered, the technical detail is here to carry the illustration, and a specialist reading closely may reach a configuration they would have set up differently. What transfers to a real engagement is the reasoning: what gets tested, how a result is rated, what evidence is required before a claim is made, and what the report refuses to claim without it. That is what this document is for.

00 DOCUMENT CONTROL

Classification, distribution, and version history

Handling

The block below is an example of the handling instructions a real engagement report carries. It does not apply to this specimen, which is fictional and published deliberately. In a real deliverable these are contractual terms rather than advice.

This report is classified **Client Confidential**. It describes exploitable weaknesses in a production environment and the precise conditions under which defensive controls failed to observe or stop an operator. It is the single most useful document an adversary could obtain about this organization.

- Store and transmit encrypted. Do not circulate over unencrypted email or unmanaged file sharing.
- Distribute on the named list below. Additions are approved by the engagement sponsor and recorded here.
- Offensive Guardian retains one encrypted copy for the retest and destroys engagement artefacts 90 days after the retest record is signed off, unless the client instructs otherwise in writing.
- Extracts may be shared with an auditor or regulator. The full attack narrative should not be.
- **No secret value appears anywhere in this document in clear text.** Not passwords, not recovered hashes, not keys or tokens, not certificate private material. Accounts are named; their secrets are not. Everything recovered during the engagement was disclosed out of band at the close-out session and is referenced here by evidence identifier only. A report is forwarded, printed, and stored for years, so it is the wrong place to put a credential even when the client already knows it.

Distribution list

NAME	ROLE	BASIS
H. Okonkwo	Chief Information Security Officer	Engagement sponsor
M. Lindqvist	Head of Cyber Defence	Trusted agent, engagement owner
R. Baptiste	SOC Manager	Trusted agent, defensive observer

NAME	ROLE	BASIS
A. Verheyden	Head of Infrastructure	Remediation owner (added 2026-05-04, approved by sponsor)
Board Risk Committee	Not applicable	Executive summary extract only (section 02)

Version history

VERSION	DATE	AUTHOR / REVIEWER	CHANGE
0.1	2026-04-28	J. Maes	Initial draft from operator log and defensive observer record
0.2	2026-05-04	J. Maes	Technical accuracy review with SOC Manager; two findings re-rated after the analyst timeline was reconciled
0.3	2026-05-11	J. Maes	Client factual-accuracy pass; ownership assigned; improvement plan sequenced with Infrastructure
0.9	2026-05-15	J. Maes	Issued as final for the operation. Retest sections held open
1.0	2026-06-19	J. Maes	Retest record added (section 12), ratings updated, residual risk statement issued

Engagement team

NAME	ROLE	INVOLVEMENT
Jean-François Maes	Founder, Offensive Guardian	Engagement lead and sole operator. Scoping, execution, evidence capture, reporting, retest.
M. Lindqvist	Head of Cyber Defence, Northwind Mutual	Client engagement owner. Objective agreement, rules of engagement, deconfliction authority.
R. Baptiste	SOC Manager, Northwind Mutual	Defensive observer. Maintained the independent defensive timeline used in section 07.

Offensive Guardian is founder-led. The person who scoped this engagement executed it and wrote this report. Where an engagement needs additional operators, they are named here before work starts.

Verification record

This report was written and reviewed by the same person. That is a real limitation of a single-practitioner firm and it is stated here rather than disguised with a second signature. What replaces a peer reviewer is a fixed verification pass, run against every issue, and recorded so a reader can see what was checked mechanically and what rests on judgement.

CHECKED	METHOD	RESULT
Every CVSS vector against its printed score	Recomputed with the v3.1 formula	All printed scores are the formula's output
Chain verdict tallies and percentages	Recounted from the matrix rows	36 behaviours, 3 not applicable, 33 assessable; summary tables match

CHECKED	METHOD	RESULT
Every cross-reference (finding, behaviour, evidence, test case, persistent issue)	Resolved mechanically	No dangling references; every referenced identifier resolves
Day numbers against calendar dates	Recomputed from day 1 = 2 March 2026	Consistent throughout
All timing intervals	Recomputed from the reconciled timeline	Consistent to the second where both records exist
The defensive timeline	Reconciled against the client's independently kept record	Two entries unreconcilable; both marked and explained
Table of contents	Re-derived from the rendered document after the technical revision	All 19 entries verified against the 84-page render

What this does not replace: an independent practitioner's judgement on whether the technical reasoning is sound. For identity and PKI-heavy engagements we recommend the client's own architect reviews the relevant findings, and we will walk through them on request at no charge. The technical content of this report was walked through with the SOC Manager on 4 May, which validates the defensive timeline but is not an independent offensive review.

: CONTENTS

What is in this report

The report is built so each audience can start where their decision is and stop when it is answered. Leadership needs sections 02 and 11. Detection engineering needs 07 through 10. Owning teams need their findings and nothing else.

How to read a finding

Every finding in section 09 carries two ratings, and they answer different questions.

RATING	ANSWERS	WHO ACTS ON IT
Path impact Critical → Informational	What the adversary achieved, or could achieve, against the business.	Risk owners deciding what to fund and in what order.
Chain verdict Held at… / Broke at…	Where the defensive chain first failed: control, telemetry, detection, or decision.	The team that owns that link. A telemetry break and a decision break need different fixes.

● HELD · THE CONTROL OR THE DECISION WORKED △ PARTIAL · WORKED IN PART ■ BROKE · THE CHAIN FAILED HERE

PIC-NN · PERSISTENT ISSUE REFERENCE

02 EXECUTIVE SUMMARY

02 The investment bought detection. It did not buy identity.

Northwind Mutual asked one question: after eighteen months and a programme that consolidated endpoint protection, migrated the SIEM, and stood up a six-person detection engineering function, is the organization measurably harder to attack? The answer is that it is harder to attack in exactly the place the programme was pointed, and not harder anywhere else.

Over eight weeks we ran an objective-led operation against the production estate under agreed rules. We reached two of three objectives. The one we did not reach was stopped cleanly, quickly, and by design: an attempt to extract policyholder records was detected in six minutes and contained in eleven. That is a real capability and it did not exist before the programme.

The other two objectives were reached, and no defensive action followed either of them. That is not the same as nobody seeing anything, and the difference is the most useful result in this report.

Four detections fired correctly during the operation, three of them on the path we took. The rule that caught our service ticket requests on day seventeen is named for exactly the technique we were using. It reached an analyst in ninety seconds and was closed four minutes later as scanner activity, which was checkable in under a minute and wrong. The same happened three more times. Your detection engineering works; the answers it produced were discarded.

Underneath that sits a second problem. The objective steps did leave records, but the records did not reach a defensive decision. The CA database retained the issued certificate request and the servicing domain controller logged a weak-mapping warning; neither source was forwarded or alerted. The disbursement API logged the state change without an approver, so a legitimate integration action and the tested action were not distinguishable by actor. The detection programme was built around endpoint telemetry, while these identity and application records sat outside its monitored paths.

The failures need different fixes and different budgets. Alert disposition needs governance, runbooks, feedback, and enough review capacity. Missing pipelines and fields need engineering.

Not one step of the compromise required an exploit. Every stage used a configuration that was already documented and already known, and several were product defaults that were never changed. This matters for how the remediation is funded: the work ahead is configuration, ownership, and logging, not a purchase.

OBJECTIVES REACHED	BLIND SPOTS YOU COULD CLOSE	ALERT ON THE FOREST OBJECTIVE	FASTEST CONTAINMENT
2 of 3	8	None	11 min
Forest control and payment authorization reached. Bulk data extraction stopped.	Of 33 assessable behaviours, 8 produced no queryable record. Eleven more were recorded but not detected.	The CA and servicing KDC retained records, but nothing generated a defensive alert.	Data egress attempt: detected at +6 min, contained at +11 min.

Objective outcomes

OBJECTIVE	AGREED WITH THE SPONSOR AS	OUTCOME	REACHED	TIMING
OBJ-1	Authorize a claims disbursement in CLARA without a legitimate second approval.	REACHED	Day 27	**No defensive containment** First defense-generated signal 4 d 06:12 after the objective. The pre-agreed reversal is recorded separately as cleanup.
OBJ-2	Demonstrate bulk extraction of policyholder records from the customer data platform.	STOPPED	Not reached	**Contained first** Time to detect 5 m 55 s, time to contain 11 m 06 s, credentials disabled at 23 m 34 s.
OBJ-3	Obtain and hold Tier-0 control of the `northwind.example` forest.	REACHED	Day 19	**No defensive alert** CA and KDC records existed, but nothing fired. Tester disclosure and relinquishment are reported separately from response metrics.

Where the defensive chain broke

Each of the 36 behaviours we ran was walked along the same chain and recorded at the point it stopped. Three of them have no defensive surface at all, so they are marked not applicable and excluded from the percentages, which are taken over the 33 assessable behaviours. This is the engagement in one table, and it is the table that should drive the budget conversation.

CHAIN VERDICT	COUNT	SHARE	WHAT IT MEANS FOR NORTHWIND
HELD AT CONTROL	8	24%	A preventive control stopped or materially changed the path. These are wins and are documented in section 10.
RAN TO A CONTAINING ACTION	2	6%	The full chain worked: logged, detected, triaged, contained. One in email, one in the data platform, and the data platform one landed before the objective.
BROKE AT TELEMETRY	8	24%	No queryable record existed, so no detection could have fired. Every one of these is collectable and is not being collected.
NOT APPLICABLE	3	Not applicable	No defensive surface exists: offline password cracking runs on our hardware, and reading a public web page is not an event Northwind can log. Recorded because they happened, excluded from the counts above.
BROKE AT DETECTION	11	33%	The activity was logged and nothing was watching for it. The data exists; forwarding, detection logic, or both are missing.
BROKE AT DECISION	4	12%	An alert fired and no containing action followed. Tooling will not fix this; triage process and runbooks will.

Read the bottom three rows by the work they require, not by size. The 11 detection breaks have source data and need forwarding, content, or both. The 8 telemetry breaks need new collection. Percentages are of the 33 assessable behaviours and are rounded. The 4 decision breaks need triage governance, runbooks, and sufficient review capacity.

The path we took

Compressed to one paragraph, because the full narrative is section 06 and it is long. A call to the IT service desk from a spoofed internal number restored an authentication factor onto an operator-controlled device, which turned a phished password into a working remote session as a claims administrator. From that desktop we requested Kerberos service tickets for the fourteen accounts carrying service principal names and recovered eleven of their passwords offline. One of the eleven, `svc-clara-int`, sat in a group holding enrollment rights on a certificate template that lets the requester choose the subject name, so we requested a certificate naming a domain administrator and the certificate authority issued it. The domain controllers were configured to accept a weakly mapped certificate, so it authenticated. That is the forest. From there, a finance file share held a hand-maintained workbook of application credentials, among them the claims platform's API integration account, and that is what turned control of the directory into the ability to move money. Twelve of the fourteen persistent-issue predicates were observed in the tested scope, one was tested and not present, and one was outside scope. Seven were load bearing in the path above.

> THE DECISION THIS ENGAGEMENT SETTLES
>
> **The programme worked. It was aimed at one plane.** Northwind now has a detection and response capability that demonstrably functions against data movement, and effectively none against identity abuse. The next increment of security spend should not buy another sensor. It should extend the discipline that already works on the endpoint to the directory, the certificate authority, and the identity recovery process.

Three moves that change the answer next quarter

These are programme-level. Host-level and configuration-level remediation is in the improvement plan (section 11), sequenced, owned, and with a retest condition attached to each item.

#	MOVE	BECAUSE	WHAT GOOD LOOKS LIKE
01	**Bring the identity plane under the detection programme.**	Identity holds the largest cluster of collectable blind spots, and more importantly every step of the path that reached two objectives ran through it.	Certificate issuance, directory object changes, and authentication factor changes are collected, retained, and carry owned detection logic on the same review cadence as endpoint logic. Service ticket requests already are; that rule works and does not need rebuilding.
02	**Turn identity recovery into a controlled process.**	A phone call restored an authentication factor with no verified check on the caller, and that single step is what put an operator inside.	Factor re-enrollment requires verification through a channel the caller did not choose, is logged as a security event, and notifies the account holder. Service desk staff are measured on refusals, not on call time.
03	**Make alert disposition auditable.**	An alert for the exact technique we used fired correctly and was closed in four minutes as scanner activity. The logic was not the problem.	Closure reasons are structured, sampled weekly, and every dismissal of a Tier-0 adjacent alert is reviewed. Detection engineering sees what happens to the alerts it ships.

What worked, and should be protected

Reported with the same evidence standard as the failures, because a report that only lists breaks gives leadership no way to tell a functioning control from an untested one. Detail in section 10.

- **Data egress detection and response.** The one objective we did not reach was stopped by a detection that fired correctly and an analyst who acted without escalation. Eleven minutes end to end.
- **The external perimeter.** Across 214 exposed services we found no unauthenticated path in and no credential that survived to a session. Initial access required a human, which is the correct shape for a perimeter.
- **Tier-0 build discipline.** The print spooler was disabled on domain controllers, DHCPv6 guard was configured on the modernized core, and SQL servers were segmented with egress restrictions. Three separate techniques died against these and are recorded as held.
- **Local administrator password management.** Managed passwords covered 96.6% of workstations. Reuse only worked inside a legacy imaging group of 43 hosts, which is a scoping gap rather than a design failure.

Retest outcome

A five-day retest in June re-ran all fourteen documented validation conditions. Of the fourteen validated breaks, nine are remediated and verified, three are partially remediated with the residual precisely stated, one is formally risk accepted with a compensating control, and one is not remediated and carries a committed date. Both demonstrated routes to Tier-0 are closed, as is the way in. The payment path is not closed: it is now detected and cancelled before money moves, which is interception rather than prevention, and it stays Critical until the September release. What remains open is named and dated rather than summarised into a single rating, for the reasons set out there. The full record, including what we could not re-verify and why, is section 12.

> Most compromises we perform during assessments have very little to do with actually exploiting many vulnerabilities. Most of the time, we rely on misconfigurations in Active Directory itself.
>
> A DECADE OF SECURITY ASSESSMENTS: SECURITY ISSUES THAT REFUSE TO DIE · SANS INSTITUTE, FEBRUARY 2025

That article is the founder's own review of ten years of assessments and is the source of the persistent issue catalogue used throughout this report (section 04.4). This engagement is a case study in its thesis: the compromise path used no novel technique and nothing that was not publicly documented years ago.

03 THE ENGAGEMENT

03 Scope, objectives, and the rules we worked under

An engagement is only as meaningful as the boundary it ran inside. This section states what was in play, what was deliberately excluded, and what constraints shaped the result, so that a reader can judge how far the evidence travels.

The decision behind the engagement

Northwind Mutual completed an eighteen-month defensive programme in late 2025: endpoint protection consolidated onto one platform, the SIEM migrated, and a detection engineering function of six people established. The Board Risk Committee asked the CISO to demonstrate the return. The engagement was commissioned to answer that question with evidence rather than coverage claims, and the objectives below were written to make the answer falsifiable.

Objectives

Objectives were agreed with the sponsor before any tooling was selected. Each is a claim the operation could prove or break, and success and failure were both defined in advance. The behaviours used to test them were selected from Northwind Mutual's own critical paths rather than from a named threat actor's playbook: this was an adversary simulation and not an emulation, for the reasons set out in section 04.

ID	OBJECTIVE	DEFINED AS ACHIEVED WHEN
OBJ-1	Authorize a claims disbursement without a legitimate second approval.	A disbursement instruction reaches the approved state in CLARA with no genuine approver action. Executed against a test policy in the production instance, agreed with Finance, capped at €1 and reversed the same day.
OBJ-2	Demonstrate bulk extraction of policyholder records.	10,000 or more records leave the POLARIS platform boundary to an operator-controlled destination. Records were to be synthetic test policies only; the extraction was blocked before any record left.
OBJ-3	Obtain and hold Tier-0 control of the forest.	Authentication as a member of Domain Admins or equivalent, held for 24 hours without loss of access.
OBJ-4	Establish where the defensive chain first breaks, for every behaviour run.	Every executed behaviour has a recorded state at control, telemetry, detection, and decision, reconciled against the defensive observer's independent timeline. This objective was met in full and produced sections 07 and 08.

OBJ-4 is the objective that distinguishes this from a penetration test. OBJ-1 to OBJ-3 ask whether an adversary can win. OBJ-4 asks what the defense does while they try, and it is answered whether the adversary wins or not.

Scope

SURFACE	IN SCOPE	NOTES
Internet-facing estate	`192.0.2.0/24`, `198.51.100.0/24` 214 responsive services	Full external assessment including authentication surfaces and exposed management interfaces.
Corporate network	`10.30.0.0/16`	User VLANs, VDI estate, print and file services.

SURFACE	IN SCOPE	NOTES
Datacentre	`10.40.0.0/16`	Application and database tiers, including CLARA and POLARIS.
Legacy segment	`10.90.0.0/16`	Retained branch infrastructure pending decommission. In scope at the client's request.
Identity	`northwind.example` forest, single domain Entra ID tenant, ADCS two-tier PKI	Full identity plane in scope. This is where the engagement was decided.
Applications	CLARA (claims disbursement), POLARIS (policyholder data), the M365 tenant, the VDI broker	Business-logic abuse in scope for CLARA and POLARIS only.
People and process	Voice-based social engineering against the IT service desk; email phishing against a 240-person sample	Individuals are never named in this report or to their management. See below.

Explicitly out of scope

EXCLUDED	REASON
The actuarial mainframe environment	Third-party operated under a separate assurance regime. Adjacency was mapped and reported; nothing was executed against it.
Managed payroll provider	Outside the client's authorization to test.
Dublin disaster recovery site	Client decision. It holds a trust with the primary forest, so the OBJ-3 result is relevant to it, but nothing was tested there and no claim is made about its state. The bounded position is in F-04.
Denial of service and availability impact	Standard exclusion. No technique was selected where degradation was a plausible side effect.
Physical intrusion, and any action against personal devices	Client decision and legal boundary respectively.
Any modification of real policyholder or claims data	Objective work used designated synthetic test records only.

Rules of engagement

- **Unannounced.** Two trusted agents knew: the Head of Cyber Defence and the SOC Manager. The SOC floor, the service desk, and infrastructure teams did not. This is what makes section 07 meaningful.
- **Deconfliction.** A dedicated number reachable within fifteen minutes throughout. Any suspected genuine incident was to be confirmed or denied by us immediately. It was invoked twice (days 11 and 33); both times we confirmed the activity was ours.
- **Abort criteria, agreed in advance.** Any impact on claims processing availability, any unintended effect on real policyholder data, any action reaching an out-of-scope system, or any instruction from the sponsor. Not invoked.
- **Attribution.** All operator traffic originated from declared ranges, logged with UTC timestamps, and is listed in Appendix A so the client can reconcile independently.
- **Credential handling.** Every credential obtained was recorded, held encrypted, disclosed in full at close-out on 24 April, and reset by the client. Nothing was retained after the retest.
- **Critical-finding disclosure.** Any finding that gives an unauthorized party control of identity, money, or bulk data is reported verbally to a trusted agent within four business hours of confirmation and in writing within twenty-four, and the operation pauses on that path until they acknowledge. The trigger is first possession of the capability, not each later use. It applied three times: the Tier-0 capable service account on day 18, forest control by certificate on day 19,

and the application credential on day 21. The day 27 disbursement used the capability already disclosed on day 21 and was not a fourth trigger. The record below shows that only the day 18 pause is fully auditable; missing notice or acknowledgement times are evidence gaps and no compliance claim is made for those rows.

- **Client revocation.** Nothing we used survived the client disabling it. At any point, without our cooperation and without notice, Northwind Mutual could have terminated every operator session by disabling the accounts involved. That is a property of how the engagement was run, not a promise.
- **Persistence.** No persistence mechanism was left beyond the operation. Tier-0 access was demonstrated for the agreed 24-hour window and relinquished at the end of it, with the enrolling certificate handed over for revocation.
- **People.** No individual who was successfully social-engineered is identified in this report, to their management, or in any artefact. The finding is about the process that permitted the action, and the report says so in the finding itself.

CRITICAL-DISCLOSURE LEDGER · SOURCE EV-32

CAPABILITY FIRST POSSESSED	POSSESSION	VERBAL NOTICE	WRITTEN NOTICE	ACKNOWLEDGEMENT, PAUSE, AND RESUME
Tier-0 capable `svc-bkp-legacy` credential	D18 09:20	D18 11:05	D18 14:20	Acknowledged in the written exchange at 14:20; path paused from 09:20 to 14:20; continuation authorized.
Certificate authentication as a domain administrator	D19 11:26	D19 13:40	Not recorded	Acknowledgement and resume times are not preserved. The report does not claim this row met the written-notice or pause requirement.
`svc-clara-api` disbursement capability	D21 time not recorded	D21 time not recorded	Not recorded	Acknowledgement and resume times are not preserved. The report does not claim this row met the timing or pause requirement.

Timeline

PHASE	DATES	CONTENT
Model	16 – 27 Feb 2026	Objective agreement, rules of engagement, threat selection, hypothesis set, defensive observer briefing.
Pressure & Observe	2 Mar – 11 Apr 2026	41 days of operation. External assessment, initial access, internal escalation, objective work.
Close-out	13 – 24 Apr 2026	Full disclosure walkthrough with the SOC, credential handover, timeline reconciliation with the defensive observer.
Report	28 Apr – 15 May 2026	Drafting, technical review, client factual accuracy pass, ownership assignment.
Improve & retest	8 – 12 Jun 2026	Re-execution of all fourteen documented validation conditions. Recorded in section 12.

04 METHOD

04 How this engagement was run and how the results are rated

A penetration test asks whether something can be broken. That question has been answered the same way for a decade, and the answer keeps being yes for the same reasons. Offensive Guardian asks a different one: while the attack is moving, what does the defensive system actually do?

WHAT KIND OF ENGAGEMENT THIS WAS

This was an adversary simulation, not an adversary emulation. The difference decides what the report can claim, so it is stated before the method rather than buried in a caveat.

Emulation starts from intelligence on a named threat actor and reproduces that actor's known tradecraft, in their known sequence, to answer "would we survive this group". *Simulation* starts from the client's own critical paths and applies realistic attacker behaviour chosen for what it will prove about the defence, without claiming to be anyone in particular.

Northwind Mutual commissioned the second thing, so no threat actor is named anywhere in this report and none should be inferred. This report can tell you how the defensive chain behaved against the behaviours listed in Appendix C. It cannot tell you how it would behave against a specific named group, because no actor profile drove the selection. Where that is the question, usually because a regulator or a board has asked it, the engagement needs an intelligence phase in front of it and is scoped separately.

04.1 The evidence loop

Every engagement runs the same four phases, regardless of size.

STEP	PHASE	WHAT HAPPENS
01	**Model**	Define the decision the organization needs to make, the threat that makes it relevant, the critical path it runs through, the safe operating constraints, and a falsifiable hypothesis for each. Written before tooling is chosen.
02	**Pressure**	Execute realistic adversary behaviour under the agreed rules, from atomic technique tests through chained objectives. Pressure is applied at the lowest level that still answers the hypothesis.
03	**Observe**	Capture the full defensive chain, not only whether the payload ran. Operator timestamps in UTC at every step, reconciled afterwards against the defensive observer's independent record. Their record is the authority on what the defense saw.
04	**Improve**	Classify the break, assign an owning team, preserve the test case so the client can re-run it without us, and prove the change through retest.

04.2 The evidence chain

Every behaviour is walked along the same six-node chain, and each node is given a state. The chain is only as strong as its weakest node: excellent detection logic behind absent telemetry is logic that never fires.

01 ACTION	02 CONTROL	03 TELEMETRY	04 DETECTION	05 DECISION	06 EVIDENCE
ADVERSARY	HELD / DEGRADED / ABSENT / N/A	COMPLETE / PARTIAL / ABSENT / N/A	FIRED / LATE / SILENT / N/A	ACTIONED / TRIAGED / MISSED / N/A	RECORDED

NODE	STATES	WHAT IS BEING JUDGED
Control	HELD / DEGRADED / ABSENT / N/A	Did a preventive or limiting control stop the action, slow it materially, or not engage at all. N/A means no client-side control surface exists.
Telemetry	COMPLETE / PARTIAL / ABSENT / N/A	Did the activity reach a queryable log with enough context to act on. *Partial* means the event exists but is missing the field that would make it actionable. N/A means the behaviour creates no client-side event.
Detection	FIRED / LATE / SILENT / SUPPRESSED / N/A	Did analytic logic produce an alert, and was it inside the window where it still mattered. *Suppressed* means logic exists and has been tuned out. N/A follows a behaviour with no defensive surface.
Decision	ACTIONED / TRIAGED / MISSED / NOT REACHED / N/A	Did a human or automation take a containing action. *Triaged* means someone looked and closed it. *Not reached* means the alert never entered a queue a person works. N/A follows a behaviour with no defensive surface.

States are always written as a label, never signalled by colour alone, so the notation survives print, greyscale, and colour-blind readers.

04.3 Rating

Findings carry two independent ratings. Collapsing them into a single severity number destroys the information the buyer needs.

Path impact · what the adversary achieved

Rated on evidence produced during this engagement, in this environment, not on a vulnerability's theoretical worst case.

CRITICAL	Achieved or trivially achievable control of the organization's identity, data, or money at scale.
HIGH	Material unauthorized access or privilege that reaches Critical with routine effort.
MODERATE	Meaningful unauthorized access, or a required link in a demonstrated chain. Load bearing in combination.
LOW	Genuine weakness with limited standalone consequence here.
INFO	Observation with no demonstrated path. Recorded for completeness.

Chain verdict · what the defense did

The furthest point the behaviour reached before it stopped. Mutually exclusive; every behaviour gets exactly one.

HELD AT CONTROL	Stopped by a preventive control. A success.
ACTIONED	Logged, detected, triaged, and contained. Where the behaviour sat on the path to an objective this is qualified *before* or *after objective*: containment that lands after the objective has prevented nothing and is never counted as a success.
NOT APPLICABLE	No client-side defensive surface exists. The behaviour is recorded but excluded from break counts and percentages.
BROKE AT TELEMETRY	Nothing recorded it. No detection was possible.
BROKE AT DETECTION	Visible, and nothing was watching.
BROKE AT DECISION	The alert fired and no containing action followed.

Reading the pair. A finding rated *Critical / Held at control* is a good day. A finding rated *Moderate / Broke at telemetry* may deserve more engineering attention, because the blind spot generalizes to every other technique that touches the same telemetry. Findings are ordered by remediation urgency, which combines both axes with effort and dependency. That ordering is a judgement and it is stated explicitly in section 11 so the client can disagree with it.

CVSS. Supplied in Appendix B where a governance process requires it. It is never the headline number, because CVSS rates a vulnerability in the abstract and this report rates what happened here.

04.4 Time to detect and time to contain

For objective-led work each stage records five defensive and outcome times, taken from two independently kept timelines. Engagement-control events are recorded separately so a required tester disclosure or planned cleanup cannot be mistaken for a defensive response.

TIMESTAMP	MEANING
t0	The adversary action.
t_recorded	The first moment any source system wrote a record of it, whether or not anyone was watching. A log line counts.
t_alert_defense	The first organization-generated defensive or operational notification put in front of a decision-maker. Scheduled tester disclosure to a trusted agent is excluded.
t_contain_defense	The first unplanned client response action that stopped or reversed the activity. Pre-agreed operator cleanup and safety reversal are excluded.
t_objective	The moment the objective was reached. Empty where the objective was prevented.
t_disclose	Tester notice to a trusted agent under the rules of engagement.
t_ack	The trusted agent's acknowledgement and authorization to resume.
t_cleanup	Pre-agreed reversal, relinquishment, or removal by the operator.

```
Time to detect = t_alert_defense − t0
Time to contain = t_contain_defense − t0
```

Whether the activity was *recorded* is reported separately from whether it was *detected*, because they call for different work. Data that exists and is never queried is a detection problem. Data that was never written is an engineering problem. The two are routinely conflated and they have different price tags.

Every populated timestamp names its source event. A field is empty where the event did not occur and says *not recorded* where the event occurred but the evidence set does not preserve a reproducible time. Disclosure and cleanup intervals remain auditable, but they never populate defensive response metrics.

Where an objective was both reached and contained, the report states which came first, because a containing action that lands after the objective has prevented nothing. That ordering is what tells a defender whether their speed was enough, and it is stated in words rather than compressed into a single derived figure.

04.5 The persistent issue catalogue

Fourteen issues account for the overwhelming majority of full compromises observed across a decade of assessments. They are catalogued so that every engagement can state plainly which were tested for, which were present, and which were not applicable.

It exists to make coverage legible. A report that finds three of these and does not say whether the other eleven were examined leaves the reader unable to tell a clean result from an untested one. The full coverage table for this engagement is section 10.4.

ID	PERSISTENT ISSUE	PRESENT AT NORTHWIND MUTUAL
PIC-01	Predictable passwords produced by complexity-and-rotation policy	TESTED: PRESENT F-09
PIC-02	Weak MFA methods, and weak recovery for strong ones	TESTED: PRESENT F-01
PIC-03	Neglected internal perimeter behind a well-maintained external one	TESTED: PRESENT structural, section 06
PIC-04	Name-resolution fallback protocols (LLMNR, NBT-NS, mDNS)	TESTED: PRESENT F-02
PIC-05	Unmanaged IPv6 permitting rogue DHCPv6 and DNS takeover	TESTED: PRESENT F-10, observed on legacy
PIC-06	Print spooler and other coerced machine authentication	TESTED: NOT PRESENT Tier-0 predicate; H-02
PIC-07	Internal web applications and admin portals with no or default authentication	TESTED: PRESENT F-07
PIC-08	Active Directory Certificate Services misconfiguration	TESTED: PRESENT F-04
PIC-09	Sensitive data exposed on shares and internal sites; inventory failure	TESTED: PRESENT F-06
PIC-10	Password reuse, especially local administrator reuse across hosts	TESTED: PRESENT F-08, observed on 43 hosts
PIC-11	Weak or never-rotated service account passwords	TESTED: PRESENT F-03
PIC-12	Default credentials on appliances and management software	TESTED: PRESENT F-07
PIC-13	SQL Server `sa` and domain-authenticated SQL abuse	NOT IN SCOPE
PIC-14	Active Directory misconfiguration chains	TESTED: PRESENT F-03, F-04

Why these issues persist

Naming the root cause matters, because the remediation for *"the same finding returns every year"* is different from the remediation for the finding itself. Every finding in section 09 states which of these five applies to it.

CAUSE	NAME	WHY THE ISSUE KEEPS COMING BACK
RC-1	**Change is expensive and risky**	Large organizations cannot easily accept downtime for a patch or the chance that a firewall rule is wrong.

CAUSE	NAME	WHY THE ISSUE KEEPS COMING BACK
RC-2	**Awareness follows testing cadence**	Organizations largely learn about these issues from assessments. Infrequent assessment means late awareness.
RC-3	**Human memory limits**	People cannot hold many long, complex, frequently rotated secrets, so they produce predictable ones.
RC-4	**Insecure defaults**	The product ships in a state that is unsafe once deployed at scale, and the hardening decision is never made. IPv6 enabled on every modern interface, the print spooler running everywhere, certificate web enrollment defaulting to HTTP and NTLM. The default is the starting condition; the missing decision is the cause.
RC-5	**Inventory failure**	Organizations cannot secure infrastructure they have not enumerated.

Source: *A Decade of Security Assessments: Security Issues That Refuse to Die*, Jean-François Maes, SANS Institute, February 2025. The catalogue is first-party: it is the founder's own published analysis of his engagement history, not a repackaged industry list.

05 LIMITS

05 What this report does not tell you

A method with no stated limits is marketing. These constraints are printed in every Offensive Guardian report and they are not boilerplate: each one changes how far a specific conclusion in this document can be carried.

It does not prove absence

The engagement proves what it tested. Behaviours we did not run are unproven, not safe. Appendix C lists every technique exercised and Appendix D every test case, so the boundary between *tested and held* and *never tested* is visible rather than implied. Where this report says a control held, it held against the specific behaviour recorded, at the time recorded.

It says nothing about any particular threat actor

This was a simulation. The behaviour set was chosen for what it would prove about Northwind Mutual's defensive chain, not because any specific group is known to use it. Nothing here supports a statement of the form "we would detect actor X", and the report should not be cited that way to a board or a regulator. That question needs an emulation with an intelligence phase, and the two are priced and scoped differently.

It is a time-boxed sample, not a census

Forty-one days of operation by one operator working under rules of engagement. A real adversary has unlimited time, no rules, no deconfliction obligation, and no requirement to avoid damage. Where a technique was excluded because it risked availability, that is noted in the finding, and the absence of the result is a limit of this engagement rather than a property of the environment.

It reflects a point in time

The estate, the detection content, and the staffing on shift all changed during the window and have changed again since. Two of the results in section 07 are directly affected by who was on shift: this is stated where it applies rather than averaged away. Evidence has a shelf life, which is the argument for validation cycles rather than an annual test.

It observes the defense as configured during the window

Three relevant conditions applied and are recorded because they affect interpretation:

- A SIEM content release on 18 March changed two detection rules mid-operation. Behaviours run before and after that date are marked in section 08.
- The SOC operated with two of eight analyst seats vacant for the whole window. The decision-node results in section 07 should be read against that staffing, not against an idealized roster.
- A planned change freeze from 30 March to 6 April suppressed some routine administrative activity, which made operator activity marginally more conspicuous, not less.

One person wrote it and one person checked it

Offensive Guardian is a single-practitioner firm. The operator, the evidence collector, the author, and the reviewer are the same person, and no amount of process fully substitutes for a second pair of eyes. The verification record in the document control section lists exactly what was recomputed mechanically, which is the part that can be made reliable without a second person. The judgement calls, which findings matter and whether the technical reasoning holds, carry the risk that comes with single-author review. Read the technical detail with that in mind, and challenge it; we would rather be corrected than believed.

Client-supplied information is taken as given

Asset inventories, network diagrams, and application ownership were provided by the client and used as supplied. Where we independently verified something, the finding says so. Where we did not, it says that too. The discrepancy between the supplied inventory and what we enumerated is itself a finding (F-12).

The retest verifies the stated condition, nothing wider

Section 12 records whether each finding's retest condition was met. It is not a re-run of the engagement. A remediated finding means that specific test case now fails for an attacker; it does not mean the underlying class of issue is eliminated everywhere in the estate. Two entries in section 12 say exactly this and explain what a wider verification would require.

> STANDING STATEMENT
>
> No engagement, and no report, guarantees that a system is secure or that an intrusion will not occur. What this document does is convert one organization's defensive claims into evidence that can be checked, argued with, and retested.

06 WHAT HAPPENED

06 Attack narrative

This is the operation in order, with the times we recorded and the findings each step produced. It is written so that a defender can follow the path and see where their own team would have had to be standing to interrupt it. All times are UTC. Day 1 is 2 March 2026.

Cross-references: `F-nn` is a validated break in section 09, `H-nn` is a control that held in section 10, `EV-nn` is an evidence item indexed in Appendix E, and `TC-nn` is a reproducible test case in Appendix D.

Days 1–8 · The perimeter did its job

We enumerated the internet-facing estate across both declared ranges and found 214 responsive services. Everything authenticating was behind the tenant's conditional access policy. Two exposed management interfaces were current on patching and neither accepted an unauthenticated request that changed state. There was no unauthenticated route in, and no credential we held at that point survived to a session. That result is recorded as `H-01` and it is a genuine one: perimeters that force an adversary through a human are perimeters that are working.

Open-source collection was more productive than the scan. A support article published on the customer-facing help portal described the internal procedure for restoring a lost authentication factor, including the three verification questions the service desk asks. It had been indexed for eleven months. The service desk's direct number is printed on the careers page. Both are reasonable individually. Together they are a script. (`F-01`, `EV-03`, root cause RC-5.)

Days 9–12 · Phishing, which failed usefully

A credential-capture campaign went to a 240-person sample drawn from public sources. 218 messages were delivered, 34 recipients opened the link, and 11 submitted credentials. **None of those credentials produced a session.** Conditional access rejected every authentication attempt from our infrastructure and the second factor was never satisfied. That is the control working as designed (`H-04`).

The campaign was also reported. The first report to the security mailbox arrived four minutes after first delivery, and the SOC had the sending domain blocked within two hours. On day 11 the SOC Manager invoked deconfliction; we confirmed the activity was ours. Both facts are positives and are recorded as such.

But a password that cannot be used is still a password. Among the eleven was the account of a claims administrator, and knowing that account's password is what made the next step credible.

PHISHING CAMPAIGN RESULT · DEFENDER-RELEVANT FIGURES EV-05

Messages sent	240	
Delivered	218	(90.8%)
Link opened	34	(15.6% of delivered)
Credentials submitted	11	(5.0% of delivered)
Sessions obtained	0	conditional access + second factor held
Reported to security mailbox	19	first report at T+4 min
Sending domain blocked		T+2 h 06 min

Day 14 · One phone call

At 14:12 we called the IT service desk from a number spoofed to an internal extension, as the claims administrator, reporting a replaced phone and a lost authenticator. The agent asked the three verification questions from the published support article. Employee number, line manager's name, and the month of joining are all recoverable from public sources and the payroll-format employee number is printed on the visitor badge photographs used in a recruitment post. We answered all three.

The agent holds the Authentication Administrator role, which the service desk tier-2 queue carries so it can resolve lockouts without escalating. That role can delete a member user's registered authentication methods and issue a Temporary Access Pass for them, and our target was an ordinary member user. The agent deleted the Authenticator registration on the user's real phone, then issued a Temporary Access Pass valid for sixty minutes and single use and read it to us. We used it to sign in and register a new Authenticator on an operator-controlled device. The call lasted six minutes and forty seconds. The Temporary Access Pass is what carried us past the registration gate that the phished password alone could not pass, because it satisfies the multi-factor requirement on the registration action itself. At 14:31 we authenticated with a password we already had and a second factor we now controlled, and at 14:44 we had a virtual desktop session as a claims administrator.

The identity provider recorded the administrator deleting the existing security information, registering the Temporary Access Pass, and the user registering new security information. The service desk also created a ticket under *Password & Access — MFA reset*. None of those records was forwarded to the SIEM or correlated with the following sign-in, and nobody notified the account holder. (`F-01`, Critical, broke at detection, `PIC-02`, root cause RC-1.)

> THE HINGE OF THE ENGAGEMENT
>
> Every technical control between the internet and a corporate desktop worked. The way in was a process that was designed to be helpful and was never designed to be a security control. This is the single step where a small change would have cost us the entire operation.

Days 15–16 · The internal perimeter

From the virtual desktop, the legacy branch segment `10.90.0.0/16` was directly reachable. It is scheduled for decommission and has been since 2024. Name-resolution fallback protocols were still enabled on that segment, so a mistyped share path from any host there is answered by whoever replies first (`F-02`, `PIC-04`, root cause RC-1).

Over eleven hours on day 16 we collected authentication material from four accounts and two machine accounts by responding to those broadcasts. Nothing in the environment recorded that we had done so. There is no telemetry for it because the segment's hosts do not forward the relevant events, which means no detection could have existed regardless of how good the detection team is.

Day 17 · The alert that fired

Authenticated as the claims administrator, we enumerated the directory and requested service tickets for the fourteen accounts carrying service principal names, in order to recover their passwords offline.

At 09:47 a SIEM rule named *Anomalous volume of Kerberos service ticket requests* fired correctly and reached the analyst queue. At 09:51 it was closed with the disposition *Vulnerability scanner — expected*. Four minutes. The detection engineering team built the right rule, the platform delivered it, and the answer was still no.

The existing analytic makes this the highest-leverage process result. The remediation is triage governance, runbooks, feedback to detection engineering, and enough review capacity to apply them. (`F-11`, broke at decision, root cause RC-2.)

Eleven of the fourteen recovered passwords were weak enough to break offline. Two of them mattered. `svc-clara-int`, whose password was set on 14 August 2019 and never changed, is a member of the group holding enrollment rights on the certificate template in the next section. `svc-bkp-legacy` is a member of a Tier-0 group through a nested group nobody had reviewed, which means that at 09:20 on day 18 we held a credential that was already sufficient for forest control (`F-03`, `PIC-11`). Possession alone triggers the critical-finding disclosure clause, so both trusted agents were told at 11:05 the same morning and acknowledged in writing at 14:20, authorizing us to continue.

We did not use it, and the reason is worth stating. Authenticating as `svc-bkp-legacy` would have met OBJ-3 a day earlier and taught the client almost nothing: one group membership change closes it. The certificate path was the more useful thing to prove, because it is available to 123 principals rather than one, and it survives the removal of any single account. The CA database retains the issued request, but no CA security-audit event was enabled or forwarded. Using the backup account first would also have contaminated the test, since we could no longer have shown that the certificate path was independently viable. Tier-0 was therefore *available* on day 18 and *demonstrated* on day 19. The report uses the later date throughout because that is when we actually authenticated. The earlier path is recorded because it is independent of the certificate path.

```
SERVICE ACCOUNT CREDENTIAL RECOVERY · SUMMARY, VALUES WITHHELD                  EV-11

SPN-bearing accounts enumerated          14
Tickets requested                        14
Recovered offline                        11    (78.6%)
Password last set, oldest                2019-08-14   svc-clara-int
Password last set, median                2021-11-02
Accounts with a password over 3 years old  9
Accounts with Tier-0 group membership      1    svc-bkp-legacy  (see F-03)

--- recovery method, so the result can be judged rather than believed ---

Ticket encryption, 11 accounts   RC4-HMAC (etype 23)      hashcat mode 13100
Ticket encryption,  3 accounts   AES256-CTS-HMAC-SHA1-96  hashcat mode 19700
First recovery at                00:04 elapsed
Eleventh recovery at             29:41 elapsed
Observed result                  11 of 11 RC4 tickets recovered
                                 0 of 3 AES256 tickets recovered

The three that survived used AES256 tickets. No conclusion about their password
strength can be drawn from this test, and none is drawn.

Reproducibility limit            The engagement log did not preserve GPU model,
                                 hashcat/runtime versions, corpus and ruleset
                                 hashes, mask charsets, phase candidate counts,
                                 or compute cost. The timings above are an
                                 observed engagement result, not a benchmark.
                                 They must not be used to predict another run.

Recovered values are withheld from this document. They were disclosed in
full at the close-out meeting on 2026-04-24 and reset by the client.
```

Days 18–19 · A certificate for the Administrator

The recovered service account is a member of `NWM-App-Service-Accounts`. That group has enroll rights on a certificate template named `NWM-Legacy-WebServer`, created in 2019 for a server migration that finished the same year. The template permits the requester to supply the subject name, includes a client authentication purpose, and does not require manager approval.

Those three properties let an enrolled principal request a certificate naming another identity. Getting one issued is not the same as being able to use it. At 11:26 on day 19 the CA issued the request and a domain controller accepted one certificate authentication as the named administrator, logging a weak-mapping warning. We held Tier-0 for the agreed twenty-four hours and met OBJ-3. The successful authentication and warning prove how the servicing KDC handled this certificate. The evidence set does not preserve that DC's OS build, last cumulative update, Audit Kerberos Authentication Service setting, or an Event 4768 result. It therefore does not attribute the March 2026 behavior to patching exception CHG-2025-0417 and does not claim Event 39 was the only certificate-bearing event available. Both trusted agents were told at 13:40. Written notice, acknowledgement, and resume times were not preserved, so this report does not claim that pause met every term of the disclosure rule. The SOC floor was not told.

The CA database retained the issued request, including request ID, disposition, requester, template, issued subject and subject alternative name, serial, and issue and expiry times. CA security auditing was not enabled, and neither the CA record nor the servicing KDC's weak-mapping warning was forwarded or alerted. One authentication was performed; this report makes no claim that the other three domain controllers processed it or logged copies. (`F-04`, Critical, broke at detection, `PIC-08`, `PIC-14`, root cause RC-4.)

CERTIFICATE TEMPLATE PROPERTIES AND BINDING STATE · NWM-LEGACY-WEBSERVER EV-14

```
Template                     NWM-Legacy-WebServer
Created                      2019-03-11
Schema version               2
Enrollee supplies subject    YES         ← requester chooses the identity
Manager approval required    NO          ← no human in the path
Authorized signatures        0
Extended key usage           Client Authentication (1.3.6.1.5.5.7.3.2)
                             Server Authentication (1.3.6.1.5.5.7.3.1)
Enroll permission            NWM-App-Service-Accounts
                             NWM-Web-Admins
Issuing CA                   NWM-ISSUING-CA-01  (audit filter: none configured)

Certificate issued to subject alternative name Administrator@northwind.example
at 2026-03-20 11:26 UTC. Serial withheld; supplied to the client for revocation.

--- why the issued certificate was accepted for authentication ---

Domain controllers                    4 x Windows Server 2022
Servicing DC                           identified in EV-26; hostname omitted here
OS build / last cumulative update      not captured
Binding registry value                 StrongCertificateBindingEnforcement = 1
Audit Kerberos Authentication Service  not captured
Event 4768 result                      not captured
SID security extension in the cert     absent
altSecurityIdentities on the target    empty
Result at the servicing KDC            authentication accepted; Event 39 logged locally
                                       and not forwarded

This evidence proves the performed authentication. It does not establish the
patch state or behavior of the three DCs that did not service the request.
```

The template properties produced the certificate and the servicing KDC accepted it. Removing the unsafe enrollment path blocks this request. Enforcing current strong certificate mapping provides a wider control, but the retained evidence does not prove which update state governed all four DCs during the test.

Days 20–26 · From the forest to the money

Forest control is not the same as business impact, and the second is what the report has to demonstrate. With directory-wide read access we searched file shares for the things that convert privilege into action.

A share named `Finance$` on `NWM-FS-03`, readable by 1,340 accounts, held a workbook called *Application Access Register*. It contained credentials for 47 applications, maintained by hand since 2021, among them `svc-clara-api`, the account the claims platform uses to authenticate its own integration calls. That is a different account from `svc-clara-int`, the internal service account recovered on day 18. The similar names are Northwind's, not ours, and the two are kept apart throughout this report because conflating them would misstate how the path was assembled. Nothing about obtaining this required domain administrator rights; a claims administrator could have read it on day 14. Possession of the claims platform credential triggered the disclosure clause again. The trusted agents were told on day 21, but the evidence set does not preserve the exact notice, written acknowledgement, or resume times, so no timing-compliance claim is made for this pause. (`F-06`, Critical, broke at telemetry, `PIC-09`, root cause RC-5.)

The same sweep found a network monitoring appliance on the datacentre segment still using the vendor's documented default credentials, with a web interface that maps the estate more accurately than the asset inventory the client supplied us (`F-07`, `PIC-07`, `PIC-12`). The discrepancy between that map and the supplied inventory is `F-12`.

Day 27 · Objective one

The claims platform enforces a second approval in its user interface. Its integration API does not. The API accepts an approval assertion from the integration service account and does not require an approver identity, because when the interface was built the only caller was a trusted internal system.

At 15:41 on day 27 we created a disbursement instruction against a designated synthetic test policy, valued at one euro by prior agreement with Finance, and moved it to the approved state using only the integration account. No human approved it. This use did not trigger a new disclosure because possession of the same capability had been disclosed on day 21. The instruction was reversed the same day and Finance confirmed reversal in writing; the exact cleanup time was not retained. (`F-05`, Critical, broke at detection.)

The first defensive signal came at 21:53 on day 31, from an overnight access review batch job that flagged the integration account's unusual API pattern. The SOC did not produce it and would not have. Nothing contained it, then or later. The first signal of any kind arrived 4 days 6 hours and 12 minutes after the objective was already met, which means an adversary finishes, banks, and is gone well before anything looks.

Day 33 · Objective two, and the day the defense won

At 10:02 we began extracting synthetic policyholder records from POLARIS toward an operator-controlled destination, targeting the agreed 10,000-record threshold.

At 10:08 two detections fired: an anomalous query volume rule owned by the detection engineering team, and a data loss prevention policy on the egress path. At 10:13 an analyst terminated the session without escalating, which was the correct call and was made on their own authority. At 10:25 the account was disabled. At 10:31 the SOC Manager invoked deconfliction and we confirmed the activity was ours.

No records left the platform boundary. OBJ-2 was not achieved. Eleven minutes from first action to containment. This is what the defensive programme bought, it is real, and it is recorded with the same rigour as the failures (`H-05`).

OBJ-2 RECONCILED TIMELINE · OPERATOR LOG AGAINST DEFENSIVE OBSERVER RECORD EV-22

```
10:02:14  operator   first bulk query issued against POLARIS
10:04:37  operator   2,140 synthetic records staged
10:08:09  defence    DETECT  "POLARIS anomalous query volume"      (det.eng owned)
10:08:41  defence    DETECT  "DLP: bulk PII egress, unmanaged dest" (platform)
10:11:02  defence    analyst opens both alerts, correlates to one session
10:13:20  defence    ACTION  session terminated
10:13:20  operator   access lost, 2,140 records staged, 0 exfiltrated
10:25:48  defence    ACTION  account disabled, password reset forced
10:31:05  defence    deconfliction call placed; activity confirmed as ours

Objective not reached: containment came first.
Time to detect  5 m 55 s
Time to contain 11 m 06 s
```

Credential ledger

Every credential the operation held, where it came from, what it unlocked, and when. Published because the one-paragraph version of this path in section 02 is the part most readers will repeat, and it should be checkable against something rather than trusted. **No secret value appears here or anywhere else in this report**; accounts are named, their secrets are not, and everything recovered was handed over at close-out and reset.

CREDENTIALS HELD DURING THE OPERATION

SOURCE	ACCOUNT	AVAILABLE	FIRST USED	WHAT IT UNLOCKED
Phishing campaign	u-claims-14	D11	D14	Nothing on its own. Conditional access refused every attempt from operator infrastructure. It mattered only because it made the day 14 phone call credible.
Service desk call	u-claims-14 second factor	D14 14:31	D14 14:44	A virtual desktop session as a claims administrator. The way in.
Name-resolution capture, legacy segment	4 user, 2 machine NetNTLMv2 material	D16	Not applicable	Nothing. Relay held on both surfaces tested (H-06, H-10). Two of the four user passwords broke offline and carried no rights worth having.

SOURCE	ACCOUNT	AVAILABLE	FIRST USED	WHAT IT UNLOCKED
Kerberos service tickets	svc-clara-int	D18 09:20	D18 13:05	Membership of `NWM-App-Service-Accounts`, which holds enrollment rights on the certificate template. The route to the forest.
Kerberos service tickets	svc-bkp-legacy	D18 09:20 disclosed 11:05	Not used	**Tier-0 group membership, by nesting.** Sufficient for forest control on day 18. Deliberately not exercised, for the reasons above.
Certificate enrollment	Administrator via issued certificate	D19 11:26	D19 11:26	Tier-0. OBJ-3. Available to any of 123 principals. The CA database retained the issued request; no CA/KDC source was forwarded or alerted.
Finance share workbook	svc-clara-api	D21 disclosed same day; exact times absent	D27 15:41	The disbursement API. OBJ-1. Readable by 1,340 accounts, so this did not require the forest at all.

WHAT THE LEDGER SHOWS THAT THE NARRATIVE DOES NOT

Two of the three objectives did not need the certificate path. Forest control was already available on day 18 through a backup service account, and the ability to move money came from a spreadsheet any of 1,340 accounts could read. The certificate path is in this report because it is the one that generalises, not because it was the only way through. Remediating it alone would have closed the most elegant route and none of the others.

Days 34–41 · Filling in the matrix

With the objectives settled, the remaining operating days were spent running the atomic behaviours needed to complete OBJ-4, including several chosen specifically because we expected them to be stopped. Coerced machine authentication via the print spooler failed against domain controllers, where the service is disabled (`H-02`). Rogue DHCPv6 advertisement was blocked on the modernized core and succeeded on the legacy segment, which is the same story as the rest of that VLAN (`F-10`). Password spraying was stopped by the second factor on every account that had one, though the detection for it fired 41 minutes after the spray finished (`F-09`). Hash coercion from a domain-joined SQL server produced material that could not be used, because those hosts cannot reach operator infrastructure (`H-03`).

The operation ended at 17:00 on day 41. Full disclosure to the SOC took place over three sessions between 13 and 24 April, including a walkthrough in which the analyst who closed the day 17 alert reviewed the decision with us. That conversation produced the strongest input into `F-11`'s remediation and it is the reason the finding is written as a process problem rather than a personal one.

07 WHAT HAPPENED

07 Defensive timeline

The operator log and the defensive observer's record were kept independently for the whole operation and reconciled after close-out. This table is the result. Where the two disagreed, the defensive record is authoritative on what the defense saw and the operator log is authoritative on what was done.

Two entries could not be reconciled to the minute because the source log had already rolled out of hot retention when we compared records on 15 April. They are marked ~ and the retention gap that caused it is itself reported as part of F-02.

RECONCILED OPERATION TIMELINE · ALL TIMES UTC

DAY / TIME	WHAT WE DID	WHAT THE DEFENSE SAW	WHAT WAS DECIDED	REF
D1–8	External enumeration, 214 services.	Edge logs recorded the traffic in full.	No alert. No review.	B-01
D9 09:00	Phishing campaign sent to 240 recipients; 218 delivered.	19 user reports; first at T+4 min. The secure email gateway held back 22 of the 240.	Domain blocked at T+2 h 06.	B-02
D9–11	11 sets of credentials submitted; authentication attempted from operator infrastructure.	Failed sign-in events with the policy reason recorded correctly.	Control blocked. No human action needed.	B-03
D11 11:20	Continued campaign infrastructure activity.	SOC correlated reports to a single actor.	Deconfliction invoked. Confirmed as ours.	None
D14 14:12	Voice call to service desk. Existing security information deleted, Temporary Access Pass issued, and a new factor registered. Call length 6 m 40 s.	Identity-provider audit records and a service desk ticket existed at source. None was forwarded to the SIEM.	No detection or account-holder notification.	B-06 · F-01
D14 14:31	Authenticated to the tenant with password plus our factor.	Sign-in recorded and successful. New device, new location, same hour as a factor change.	No rule correlates those three facts.	B-07 · F-01
D14 14:44	Virtual desktop session as a claims administrator.	Session recorded as normal use.	No action	B-07
D15 ~10:00	Legacy segment reachability confirmed from VDI.	No flow data on that segment.	Not reached	B-08
D16 08:10 –19:20	Name-resolution poisoning on the legacy segment for 11 h. Material collected from 4 user and 2 machine accounts.	Nothing forwarded. Host logs existed locally and had rolled by the time we compared.	Not reached	B-09 · F-02
D16 ~20:00	Relay attempted against 61 SMB hosts, then against the certificate web enrollment endpoint over HTTP.	Signing negotiation failures on the SMB targets; the enrollment endpoint rejected the HTTP request, but the retained IIS evidence is insufficient to identify the cause.	Both attempts failed. No alert on either.	B-11 · B-36 H-06 · H-10

DAY / TIME	WHAT WE DID	WHAT THE DEFENSE SAW	WHAT WAS DECIDED	REF
D17 09:44	Directory enumeration over LDAP.	Query content not logged.	No alert	B-12
D17 09:46	Service tickets requested for 14 accounts.	**Alert fired at 09:47**, rule *Anomalous volume of Kerberos service ticket requests*. Reached the analyst queue.	**Closed 09:51** as *Vulnerability scanner — expected*. 4 minutes.	B-13 · F-11
D17 10:00 –D18	Offline password recovery. 11 of 14 recovered.	Not observable by design.	Not applicable	B-14 · F-03
D18 13:05	Certificate template enumeration.	No issuance or enrollment auditing on the CA.	Not reached	B-17 · F-04
D19 11:26	**Certificate issued naming a domain administrator. OBJ-3 reached.**	CA database record retained. The servicing DC logged one weak-mapping warning. Neither source was forwarded.	No defense-generated alert. Tester disclosure at 13:40; written acknowledgement time not recorded.	B-18/19 · F-04
D20 09:15	Directory replication requested from a non-domain-controller source.	Alert fired at 09:16: *Replication request from non-DC source*.	Closed 09:58 as *backup software*. 42 minutes.	B-22 · F-11
D21–26	Share enumeration; credential register located on `NWM-FS-03`.	Object access auditing not enabled on that share.	Not reached	B-23/24 · F-06
D24 16:40	Lateral authentication using a reused local administrator credential across 43 hosts.	Alert fired: *Local admin authentication from unexpected source*.	Closed as *imaging activity*. No verification with the desktop team.	B-16 · F-08
D27 15:41	**Disbursement approved via API with no approver. OBJ-1 reached.** €1, synthetic policy.	API call logged without approver identity, so the record is indistinguishable from a legitimate integration call.	No defensive action. Pre-agreed reversal completed the same day; exact cleanup time not retained.	B-31 · F-05
D31 21:53	No operator action	Overnight access review job flags the integration account's API pattern.	Ticket raised to application owner. Not treated as security.	F-05
D33 10:02	Bulk extraction of synthetic records begins.	**Two detections at 10:08**: anomalous query volume and DLP bulk egress.	**Session terminated 10:13. Account disabled 10:25.**	B-30 · H-05
D33 10:31	Confirmed activity as ours.	SOC Manager invoked deconfliction.	Correct process, correctly followed.	None
D34–41	Remaining atomic behaviours run to complete the matrix, including several expected to be stopped.	Mixed. See section 08 row by row.	Mixed.	section 08
D41 17:00	Operation ends. Access relinquished.	Not applicable	Not applicable	None

Time to detect, time to contain

Read against section 04.4. *Defense alerted* means an organization-generated notification reached a normal decision path. Scheduled tester disclosure and pre-agreed cleanup are excluded from the two response metrics and printed separately below. Every populated time names its evidence source.

OBJECTIVE	ACTION, T0	OBJECTIVE	RECORDED	DEFENSE ALERTED	DEFENSE CONTAINED
OBJ-1	D27 15:41 API approval, EV-19	D27 15:41 state approved, EV-19	D27 15:41 CLARA state-transition log, EV-19	D31 21:53 +4 d 06:12, access-review ticket, EV-20	None
OBJ-2	D33 10:02:14 first bulk query, EV-22	Not reached	D33 10:02:14 first query record, EV-22	D33 10:08:09 +5 m 55 s, query-volume alert, EV-22	D33 10:13:20 +11 m 06 s, session termination, EV-22
OBJ-3	D19 11:26 certificate request and authentication, EV-14/26	D19 11:26 Tier-0 authentication, EV-26	D19 11:26 CA database record, EV-15; servicing-DC Event 39, EV-26	None	None

ENGAGEMENT-CONTROL EVENTS EXCLUDED FROM DEFENSIVE-RESPONSE METRICS

OBJECTIVE	T_DISCLOSE	T_ACK	T_CLEANUP
OBJ-1	D21, exact time not recorded. First possession of `svc-clara-api`, EV-32.	Not recorded, EV-32.	D27, same day; exact time not recorded. Finance reversal confirmation, EV-20.
OBJ-2	Not applicable. The defense contained the test and deconfliction followed.	Not applicable.	Not applicable.
OBJ-3	D19 13:40, tester notice to both trusted agents, EV-32.	Not recorded, EV-32.	After the agreed 24-hour hold; exact time not separately recorded. Certificate transferred for revocation, EV-15.

WHAT THE THREE COLUMNS SAY TOGETHER

Northwind's response capability was fast when engaged. Time to contain on OBJ-2 was eleven minutes, from two independent detections and an analyst acting on their own authority. OBJ-1 and OBJ-3 both left source records, but neither generated a defensive alert in time to produce containment. The investment question is therefore collection quality, alerting, and ownership around identity and application activity.

08 WHAT HAPPENED

08 Evidence matrix

Every behaviour executed during the engagement, walked along the chain and recorded at the point it stopped. This is the working document for detection engineering: each row is a test case that can be re-run, and the state columns say precisely which link needs work.

CONTROL HELD / DEGRADED / ABSENT TELEMETRY COMPLETE / PARTIAL / ABSENT DETECTION FIRED / LATE / SILENT

DECISION ACTIONED / MISSED / NOT REACHED N/A NO DEFENSIVE SURFACE EXISTS

† AFFECTED BY THE 18 MAR CONTENT RELEASE

How PARTIAL telemetry is judged. A partial record breaks the chain at telemetry when the missing field is the one a detection would have to key on, and at detection when a rule could have been written against what exists. The judgement is stated in the relevant finding so it can be disputed.

36 VALIDATED BEHAVIOURS · FULL MATRIX

ID	BEHAVIOUR	ATT&CK	CONTROL	TELEMETRY	DETECTION	DECISION	CHAIN VERDICT	REF
B-01	External service enumeration across both declared ranges	T1595.002	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	H-01
B-02	Credential phishing from a lookalike domain	T1566.002	DEGRADED	COMPLETE	FIRED	ACTIONED	Actioned	H-08
B-03	Authentication with phished credentials from operator infrastructure	T1078.004	HELD	COMPLETE	NOT REACHED	NOT REACHED	Held at control	H-04
B-04	Password spray against the tenant, 1,180 accounts	T1110.003	HELD	COMPLETE	LATE	NOT REACHED	Held at control	H-04 · F-09
B-05	Collection of internal process documentation from public sources	T1593	N/A	N/A	N/A	N/A	Not applicable	F-01
B-06	Voice social engineering: authentication factor re-enrollment	T1556.006	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-01
B-07	Authentication with the re-enrolled factor from a new device and location	T1078.004	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-01
B-08	Reachability test from the VDI estate to the legacy segment	T1046	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-02
B-09	Name-resolution response poisoning (LLMNR, NBT-NS)	T1557.001	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-02
B-10	Offline recovery of captured network authentication material	T1110.002	N/A	N/A	N/A	N/A	Not applicable	F-09
B-11	NTLM relay to SMB services across 61 internal hosts	T1557.001	HELD	COMPLETE	SILENT	NOT REACHED	Held at control	H-06
B-12	Authenticated directory enumeration over LDAP	T1087.002	ABSENT	PARTIAL	SILENT	NOT REACHED	Broke at detection	F-12
B-13	Bulk service ticket request for SPN-bearing accounts	T1558.003	ABSENT	COMPLETE	FIRED	MISSED	Broke at decision	F-11 · F-03
B-14	Offline password recovery from service tickets	T1110.002	N/A	N/A	N/A	N/A	Not applicable	F-03
B-15	AS-REP roasting against accounts without pre-authentication	T1558.004	HELD	COMPLETE	NOT REACHED	NOT REACHED	Held at control	H-07

ID	BEHAVIOUR	ATT&CK	CONTROL	TELEMETRY	DETECTION	DECISION	CHAIN VERDICT	REF
B-16 †	Lateral authentication with a reused local administrator credential	T1078.003	DEGRADED	COMPLETE	FIRED	MISSED	Broke at decision	F-08 · F-11
B-17	Certificate template and enrollment permission enumeration	T1649	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-04
B-18	Certificate request with a requester-supplied subject name	T1649	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-04
B-19	Certificate-based authentication as a domain administrator	T1550	ABSENT	PARTIAL	SILENT	NOT REACHED	Broke at detection	F-04
B-20	Certificate web enrollment reachable over HTTP with NTLM	T1649	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-04
B-21	Coerced machine authentication via the print spooler against domain controllers	T1187	HELD	COMPLETE	SILENT	NOT REACHED	Held at control	H-02
B-22 †	Directory replication requested from a non-domain-controller source	T1003.006	ABSENT	COMPLETE	FIRED	MISSED	Broke at decision	F-11

36 VALIDATED BEHAVIOURS · FULL MATRIX (CONTINUED)

ID	BEHAVIOUR	ATT&CK	CONTROL	TELEMETRY	DETECTION	DECISION	CHAIN VERDICT	REF
B-23	Estate-wide file share enumeration	T1135	ABSENT	PARTIAL	SILENT	NOT REACHED	Broke at detection	F-06
B-24	Credential harvesting from share content	T1552.001	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-06
B-25	Access to a monitoring appliance with vendor default credentials	T1078.001	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-07
B-26	Rogue DHCPv6 server replies (ADVERTISE/REPLY) on the modernized core	T1557.003	HELD	COMPLETE	SILENT	NOT REACHED	Held at control	H-09
B-27	Rogue DHCPv6 server replies and router advertisements, legacy segment	T1557.003	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-10
B-28	Hash coercion from a domain-joined database server	T1187	HELD	PARTIAL	SILENT	NOT REACHED	Held at control	H-03
B-29	Scheduled task persistence on a member server	T1053.005	DEGRADED	COMPLETE	FIRED	MISSED	Broke at decision	F-11
B-30	Bulk query and staging of records from the data platform	T1213	DEGRADED	COMPLETE	FIRED	ACTIONED	Actioned before objective	H-05
B-31	Approval of a disbursement through the integration API without an approver	No mapping	ABSENT	PARTIAL	SILENT	NOT REACHED	Broke at detection	F-05
B-32	Inbox rule creation to hide matching messages in the same mailbox	No mapping	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-13
B-33	Application consent grant and use of the resulting access token	T1671 T1550.001	ABSENT	COMPLETE	SILENT	NOT REACHED	Broke at detection	F-14
B-34	Group Policy object modification from a Tier-0 context	T1484.001	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-04
B-35	Domain trust enumeration toward the disaster recovery forest	T1482	ABSENT	ABSENT	NOT REACHED	NOT REACHED	Broke at telemetry	F-04
B-36	NTLM relay to certificate web enrollment over HTTP	T1557.001	HELD	COMPLETE	SILENT	NOT REACHED	Held at control	H-10

Where the gaps cluster

The matrix is more useful grouped than listed. Every behaviour is assigned to one plane and the identifiers are printed, so this table can be checked against the matrix above rather than taken on trust.

Three behaviours have no defensive surface and are excluded. Offline password cracking happens on our hardware (`B-10` , `B-14`) and reading a page on your public help portal is not an event you can log (`B-05`). They carry the verdict *not applicable* in the matrix and are excluded from every count and percentage here, because counting them against Northwind would inflate the number and misdirect the spend.

ALL 36 BEHAVIOURS BY PLANE, WITH THE BLIND SPOTS IDENTIFIED

PLANE	BEHAVIOURS	COLLECTABLE BLIND SPOTS	WHICH ONES, AND WHAT IT MEANS
Identity and directory	17	3	Members: `B-03` , `B-04` , `B-06` , `B-07` , `B-10` , `B-12` - `B-15` , `B-17` - `B-20` , `B-22` , `B-34` - `B-36` . Blind spots: `B-17` , `B-34` , `B-35` . `B-06` and `B-18` had source records and are detection breaks; `B-10` and `B-14` are offline and not applicable.

PLANE	BEHAVIOURS	COLLECTABLE BLIND SPOTS	WHICH ONES, AND WHAT IT MEANS
Network	7	3	Members: `B-01`, `B-08`, `B-09`, `B-11`, `B-26` – `B-28`. Blind spots: `B-08`, `B-09`, `B-27`, all on the legacy segment.
Application and data	5	2	Members: `B-23` – `B-25`, `B-30`, `B-31`. Blind spots: `B-24`, reading the credential workbook, and `B-25`, appliance administration. `B-30` ran the full chain to containment before its objective.
Endpoint	3	0	Members: `B-16`, `B-21`, `B-29`. Every row produced usable data; the failures occurred at decision.
Email and tenant	3	0	Members: `B-02`, `B-32`, `B-33`. All were collected; `B-02` ran the full chain to a containing action.
Outside the estate	1	None	`B-05`, collecting a published support article. Verdict *not applicable*. Not a logging gap, and the fix is editorial rather than technical.

36 behaviours, of which 3 are not applicable and 33 are assessable. 8 of those 33 produced no queryable record, which is the telemetry-break figure in the summary table.

THE FINDING UNDER THE FINDINGS

The eight telemetry blind spots cluster in three places: three identity and directory behaviors, three legacy-network behaviors, and two application and data behaviors. The legacy segment already has a dated decommission commitment. The other gaps need collection or source onboarding.

Counts alone do not decide the next budget. The factor-change and certificate-issuance records existed but were not monitored, while directory changes and trust enumeration produced no queryable record. Together those telemetry and detection gaps covered the path to forest control. The work is therefore a mix of source onboarding and detection content, not a claim that identity activity left no records.

09 RESULTS

09 Validated breaks

Fourteen findings, ordered by remediation urgency rather than by either rating alone. Urgency combines path impact, where the chain broke, the effort involved, and what depends on what. The ordering is a judgement and it is set out below so it can be argued with.

FINDING REGISTER

STABLE ID / ALIAS	FINDING	PATH IMPACT	CHAIN VERDICT	PERSISTENT ISSUE	RETEST
OG-2026-0143-01 F-01	Authentication factors are restored by telephone with no verified check	CRITICAL	Broke at detection	PIC-02	REMEDIATED
OG-2026-0143-04 F-04	A certificate template let an enrollee authenticate as another identity, and its records were not monitored	CRITICAL	Broke at detection	PIC-08·14	REMEDIATED
OG-2026-0143-06 F-06	An application credential register is readable by 1,340 accounts	CRITICAL	Broke at telemetry	PIC-09	REMEDIATED
OG-2026-0143-03 F-03	Service account passwords have not changed since 2019, and one holds Tier-0 rights	CRITICAL	Not applicable	PIC-11	PARTIAL
OG-2026-0143-05 F-05	The disbursement API approves payments without an approver identity	CRITICAL	Broke at detection	Not scored	PARTIAL
OG-2026-0143-11 F-11	Tier-0 adjacent alerts are closed without verification	HIGH	Broke at decision	Not scored	REMEDIATED
OG-2026-0143-02 F-02	Name-resolution fallback protocols are enabled on a segment nothing logs	HIGH	Broke at telemetry	PIC-04	REMEDIATED
OG-2026-0143-07 F-07	A monitoring appliance still uses vendor default credentials	MODERATE	Broke at telemetry	PIC-07·12	REMEDIATED
OG-2026-0143-14 F-14	Users can grant an application ongoing mail and file access without review	HIGH	Broke at detection	Not scored	REMEDIATED
OG-2026-0143-08 F-08	Local administrator credentials are reused across 43 legacy hosts	HIGH	Broke at decision	PIC-10	REMEDIATED
OG-2026-0143-13 F-13	Inbox rules can hide approval mail, and nothing watches them	MODERATE	Broke at detection	Not scored	REMEDIATED
OG-2026-0143-09 F-09	The password policy reliably produces predictable secrets	MODERATE	Held at control	PIC-01	PARTIAL

STABLE ID / ALIAS	FINDING	PATH IMPACT	CHAIN VERDICT	PERSISTENT ISSUE	RETEST
OG-2026-0143-10 F-10	IPv6 is unmanaged on the legacy segment	MODERATE	Broke at telemetry	PIC-05	RISK ACCEPTED
OG-2026-0143-12 F-12	The asset inventory does not describe the estate	LOW	Broke at detection	PIC-09	NOT REMEDIATED

On the ordering. F-11 sits above three other High findings because it is the only one whose fix costs nothing and whose failure invalidates work already paid for: correct detections were built, shipped, fired, and dismissed. F-12 sits last despite being the root cause of F-07, F-08, F-13 and F-14, because it is a programme rather than a fix and the interim mitigation is already in the plan. F-14 sits above the other Highs because it granted ongoing delegated access to mail and files, while its preventive and detection fixes are bounded tenant changes. F-05 sits below three Criticals with lower business impact because its remediation depends on a release train that cannot be accelerated, and the compensating control in F-06 removes most of the exposure sooner.

OG-2026-0143-01 display alias F-01

Authentication factors are restored by telephone with no verified check

CRITICAL PATH IMPACT | BROKE AT DETECTION | PIC-02 | RC-1 | T1556.006 | T1598.004

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
VISHING	ABSENT	COMPLETE	SILENT	NOT REACHED	EV-03 · 07

SCOPE

The IT service desk factor-reset procedure. The tier-2 queue holds Authentication Administrator, which covers the 4,180 member accounts. Privileged accounts are handled by the identity team, who hold Privileged Authentication Administrator, **using the same three verification questions**. We tested the member path; the privileged path was not exercised, and the finding treats it as equally exposed because the verification step is identical, which the client accepted at close-out. Primary behaviour `B-06`. Supporting behaviours `B-05`, `B-07`. Distribution: 2 broke at detection, 1 not applicable.

ADVERSARY ACTION

We called the service desk from a number spoofed to an internal extension, as a claims administrator whose password we already held, reporting a replaced phone. The agent asked the three verification questions specified in the internal procedure. All three answers (employee number, line manager, month of joining) were recoverable from public sources, and the procedure itself was published on the customer-facing help portal and indexed for eleven months. The agent, holding Authentication Administrator, deleted the existing Authenticator registration and issued a Temporary Access Pass, single use and valid for sixty minutes, reading it out over the phone. We used it to register a new Authenticator on our own device. Six minutes and forty seconds, start to finish.

CONTROL

ABSENT. Knowledge-based verification is not a control when the knowledge is public. There is no callback to a number of record and no manager confirmation. The privileged path sits with a different team holding a higher role, but the verification questions it asks are the same three, so the additional role separates who can act rather than how well the caller is checked.

TELEMETRY

COMPLETE AT SOURCE. The identity provider recorded the administrator deleting the existing security information, registering the Temporary Access Pass, and the user registering new security information. The records identify actor, target, activity, result, and time and are queryable during the 30-day retention window. The service desk ticket records the authorizing request. None was forwarded to the SIEM. Under this report's source-queryable definition, forwarding is a pipeline gap, not absent telemetry.

DETECTION	`SILENT.` No rule consumed the source events or correlated deletion, TAP creation, method registration, and the following unfamiliar-device sign-in.
DECISION	`NOT REACHED.` The account holder was never notified, so the one person certain to know the request was fraudulent was never asked.
CONSEQUENCE	Multi-factor authentication is only as strong as its recovery path, and this recovery path is a phone call. Any account in the organization, including Tier-0 administrators, can be taken over by an adversary who holds a password and can read LinkedIn. This finding is what put an operator inside; every subsequent finding in this report was reached through it.
ROOT CAUSE	**RC-1.** The procedure was written for availability, and tightening it adds handling time to a team measured on handling time. Contributing: **RC-5**, in that nobody owned the question of what the public help portal was publishing.
OWNER	IT Service Management (procedure) · Identity & Access (technical controls and logging) · Corporate Communications (published content)
REMEDIATION	1. **Verify through a channel the caller did not choose.** Call back on the number of record, or require confirmation from the line manager inside the ticket workflow. This single change would have cost us the operation. 2. **Prohibit service desk factor reset for privileged and Tier-0 adjacent accounts entirely.** Those go through a two-person process with identity verification in person or on video. 3. **Forward the identity provider's audit events for security-information deletion, Temporary Access Pass issuance, and authentication method registration**, and alert on the sequence: a method deleted, a pass issued, a method registered, and a sign-in from an unfamiliar device inside 24 hours. The source events exist; the pipeline and correlation are missing. 4. **Restrict who can issue a Temporary Access Pass and to whom.** Scope the policy so the help desk cannot issue one for a privileged account at all, and require a second approver for anyone else. 5. **Notify the account holder out of band on every factor change**, to an address or number the change did not modify. 6. **Remove the verification procedure from public content** and treat verification scripts as internal-only, reviewed annually. 7. **Measure the service desk on correct refusals, not handle time**, and run authorized simulated calls quarterly. Staff cannot be asked to be a control and then penalized for behaving like one. Trade-off to state plainly: items 1 and 2 add time to a genuine lockout and will produce complaints. The alternative is the path in this report.
RETEST	`TC-01` · An authorized simulated call using only publicly available information about the target does not result in a factor change; *and* any factor change performed by the service desk appears in the SIEM within 15 minutes and is correlated with any subsequent sign-in from an unfamiliar device.
REFERENCES	NIST SP 800-63B, sections 5.1.3.3 and 6.1.2.3 (authenticator binding and recovery) · MITRE ATT&CK T1556.006, T1598.004 · CISA guidance on help desk social engineering

OG-2026-0143-04 display alias F-04

A certificate template let an enrollee authenticate as another identity, and its records were not monitored

CRITICAL PATH IMPACT | BROKE AT DETECTION | PIC-08 | PIC-14 | RC-4 | T1649 | T1550

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
CERT REQUEST	ABSENT	COMPLETE	SILENT	NOT REACHED	EV-14 · 15

SCOPE

Template `NWM-Legacy-WebServer` on `NWM-ISSUING-CA-01`; enroll rights held by `NWM-App-Service-Accounts` (114 principals) and `NWM-Web-Admins` (9 principals). The Dublin recovery forest was out of scope and not tested. B-35 enumerated the trust and nothing further, so what can be said is bounded: a principal in this forest could authenticate across the trust, subject to the target forest's own permissions. Whether that yields privilege there depends on trust direction, selective authentication, SID filtering, and resource ACLs, none of which were examined. This report makes no claim that the recovery forest is compromised. Primary behaviour `B-18`. Supporting behaviours `B-17`, `B-19`, `B-20`, `B-34`, `B-35`. Distribution: 3 broke at telemetry, 3 broke at detection.

ADVERSARY ACTION

Holding a recovered service account in one of the two enrolled groups, we requested a certificate specifying a subject alternative name of `Administrator@northwind.example`. The certificate authority issued it. We authenticated as that identity and held Tier-0 for the agreed twenty-four hours, meeting OBJ-3 on day 19.

CONTROL

`ABSENT.` Three template properties combine into an enrollment path: the requester supplies the subject name, the template carries a client authentication purpose, and manager approval is not required. Removing any one of the three blocks the tested request. The template was created in March 2019 for a server migration that completed the same year.

Issuance is only half of the demonstrated path. The servicing KDC accepted the issued certificate and logged Event 39. That observation proves the performed authentication. The evidence set does not retain the DC build, last cumulative update, or complete Kerberos audit-policy state, so this report does not attribute the result to patching exception CHG-2025-0417 or generalize it to all four domain controllers.

TELEMETRY

`COMPLETE AT ISSUANCE FOR THE PRIMARY BEHAVIOUR.` The CA database retained request ID, disposition, requester, template, issued subject and subject alternative name, serial, and issue and expiry times. CA security auditing was disabled, and the CA record was not forwarded.

`PARTIAL AT AUTHENTICATION.` The servicing domain controller logged Event 39 locally and it was not forwarded. Only one certificate authentication was performed. Audit Kerberos Authentication Service and Event 4768 were not captured, so no uniqueness claim is made for Event 39 and no event is attributed to the three DCs that did not service the request.

DETECTION

`SILENT.` No analytic queried the CA database or consumed the servicing KDC record. The evidence supports an unmonitored-record finding, not an absence-of-record finding.

DECISION	`NOT REACHED.` No defense-generated alert entered an operational queue during the 41-day window. Tester disclosure to trusted agents is recorded separately under the rules of engagement.
CONSEQUENCE	Each of the 123 enrolled principals could request a certificate carrying an arbitrary subject under this template. The performed request was issued and one KDC accepted it as a domain administrator, meeting OBJ-3. Historical issuance can be queried in the CA database, but the report did not complete a lifetime abuse review and cannot determine intent from issuance records alone.
ROOT CAUSE	**RC-5.** The migration template remained published seven years after its purpose ended and had no owner or retirement review. Contributing: **RC-1**, because the legacy certificate-authentication dependency had no dated modernization decision. The report does not assign the demonstrated KDC behavior to an unverified patch state.
OWNER	Identity & Access (PKI) · Security Engineering (CA audit forwarding and detection content)
REMEDIATION	1. **Remove the requester-supplied-subject property from the template, or delete the template.** The migration it served finished in 2019. This is a same-day change and it closes the path. 2. **Require manager approval on any template carrying a client authentication purpose where the subject can be supplied.** Applies to three further templates found during the sweep, which have narrower enroll permissions and were not exploited. 3. **Restrict enroll permissions to named, minimal groups.** A 114-member service account group should not hold enroll rights on an authentication-capable template. 4. **Enable issuance auditing on the certificate authority** for issued certificates, failed requests, and template changes, and forward to the SIEM. Without this, none of the detection work below is possible. 5. **Bring every domain controller to a supported cumulative-update level and verify full strong certificate mapping in a directed test against each DC.** Record OS build, last update, binding configuration, and the result before retiring CHG-2025-0417. 6. **Collect certificate-related Kerberos events from every domain controller**, including Event 39 and certificate fields in Event 4768 where the audit policy produces them. Alert on weak mapping and privileged certificate authentication. 7. **Revoke the certificate issued during this engagement** (serial supplied at close-out) and review the CA database for the template's lifetime. Treat intent as unresolved where issuance records alone cannot establish it. 8. **Build detection content for certificate-based authentication**: a certificate issued with a subject that does not match the requesting principal, and any authentication using a certificate issued in the last hour.
RETEST	`TC-04` · A principal in `NWM-App-Service-Accounts` cannot obtain a certificate naming another identity from any template in the forest. Any certificate issued produces an audit event in the SIEM within 15 minutes, and any certificate-based authentication is recorded with the issuing template and serial number.
REFERENCES	SpecterOps, *Certified Pre-Owned: Abusing Active Directory Certificate Services*, 2021 · Microsoft KB5014754, certificate-based authentication changes · MITRE ATT&CK T1649 · CISA/NSA/NCSC joint guidance on detecting and mitigating Active Directory compromises

OG-2026-0143-06 display alias F-06

An application credential register is readable by 1,340 accounts

CRITICAL PATH IMPACT | BROKE AT TELEMETRY | PIC-09 | RC-5 | T1552.001

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
SHARE READ	ABSENT	ABSENT	NOT REACHED	NOT REACHED	EV-18

SCOPE	`\\NWM-FS-03\Finance$`, a workbook titled *Application Access Register* containing credentials for 47 applications, maintained by hand since 2021. Share permissions resolve to 1,340 accounts. Primary behaviour `B-24`; supporting `B-23`. Distribution: 1 broke at telemetry, 1 broke at detection.
ADVERSARY ACTION	We enumerated shares across the estate and read the file. No privilege beyond an ordinary domain account was required at any point. The claims administrator account obtained on day 14 could have read this on day 14; we found it on day 21 only because we looked in a different order.
CONTROL	ABSENT. The share access control list grants read to a broad group inherited from a 2021 reorganization. The file is not encrypted, not rights-managed, and not classified.
TELEMETRY	ABSENT. Object access auditing is not enabled on the share, so there is no record of who has read this file, at any point in its five-year existence. Share enumeration itself is *partially* visible through session logs.
DETECTION	SILENT. A rule against bulk share enumeration could have been built on what exists and was not.
DECISION	NOT REACHED.
CONSEQUENCE	Forty-seven applications are effectively credentialed to 1,340 people. One of them is the claims platform integration account, which is what converts this finding into the ability to move money (F-05). The combined exposure is not "an adversary with domain administrator rights can approve a payment", it is "any of 1,340 accounts can". **This should be treated as a credential compromise event, not only as a finding.**
ROOT CAUSE	**RC-5.** The file exists because a real need exists and no capability was provided for it. Removing the file without providing that capability guarantees it returns under a different name in a different folder. Contributing: **RC-1**, in that the share permissions have not been revisited since the reorganization that widened them.
OWNER	Identity & Access (secret management capability) · Infrastructure (share permissions and auditing) · Finance (data owner)

REMEDIATION

1. **Treat all 47 credentials as compromised and rotate them.** Not "review", rotate. The read history is unknowable.
2. **Provide a secrets manager with delegated access before removing the file**, and migrate the register into it. This is the item that decides whether the finding recurs.
3. **Enable object access auditing on shares holding credential or regulated data** and alert on bulk read.
4. **Run a credential-pattern content scan across all file shares.** The confirmed workbook warrants a bounded search for the same pattern elsewhere.
5. **Reduce the `Finance$` access control list to the population that needs it**, and review inherited permissions from the 2021 reorganization more widely.

RETEST

`TC-06` · The workbook and any copy is absent from all indexed shares; all 47 credentials show a set date after the remediation date; a credential-pattern scan returns no hits above the agreed threshold; and a bulk read of `Finance$` generates an alert within 15 minutes.

REFERENCES

MITRE ATT&CK T1552.001, T1135 · CIS Critical Security Control 3 (Data Protection) · CWE-522, Insufficiently Protected Credentials

OG-2026-0143-03 display alias F-03

Service account passwords have not changed since 2019, and one holds Tier-0 rights

CRITICAL PATH IMPACT | NOT APPLICABLE | PIC-11 | RC-3 | T1558.003 | T1110.002

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
OFFLINE RECOVERY	N/A	N/A	N/A	N/A	EV-11

SCOPE

14 accounts carrying service principal names in `northwind.example`. 11 passwords recovered offline. 9 accounts have passwords older than three years. One, `svc-bkp-legacy`, is a member of a Tier-0 group. Three of the 14 principal names point at hosts that no longer exist. Primary behaviour `B-14`; supporting `B-13`. Distribution: 1 not applicable, 1 broke at decision.

ADVERSARY ACTION

As an ordinary authenticated domain user we requested service tickets for all 14 accounts and recovered 11 passwords offline. Eleven tickets used RC4-HMAC and all 11 were recovered; three used AES256 and were not recovered. That result does not establish the strength of the three unrecovered passwords. The run reached its eleventh recovery at 29 hours 41 minutes, but EV-11 does not retain the hardware model, software versions, candidate-set hashes, mask spaces, candidate counts, or compute cost. The timing is therefore an observed result, not a reproducible benchmark or a basis for predicting another run.

CONTROL

ABSENT. No managed service accounts are in use. The domain password policy applies a 12-character minimum that these accounts predate, and no rotation has been enforced on them. The oldest password was set on 14 August 2019.

TELEMETRY

NOT APPLICABLE to primary behaviour B-14. Offline recovery occurs on adversary infrastructure and creates no client-side event. Treating that modeling boundary as missing telemetry would charge the client for an impossible sensor. The ticket request in supporting behaviour B-13 was logged and detected; its missed decision is F-11.

DETECTION

NOT APPLICABLE to B-14. See F-11 for B-13.

DECISION

NOT APPLICABLE to B-14.

CONSEQUENCE

The test recovered 11 service-account passwords, including a credential with Tier-0 group membership. That is a second, independent route to forest compromise that does not touch F-04. The result proves exposure in this engagement; the incomplete compute record does not support a general cracking-time claim.

ROOT CAUSE

RC-3. The passwords were chosen by people, for accounts nobody expected to have to remember. Contributing: **RC-1**. Rotating a service account password risks an outage in a system whose dependencies are not documented, so it is deferred indefinitely.

OWNER

Identity & Access · the application owner of each dependent service

REMEDIATION

1. **Convert to group managed service accounts wherever the application supports it.** Rotation becomes automatic and the password becomes unrecoverable at this scale. This is the fix; everything below is for the exceptions.
2. **Remove `svc-bkp-legacy` from the Tier-0 group** and grant the specific rights the backup job requires. Do this first: it is a same-day change and it removes the worst outcome.
3. **Remove the three orphaned service principal names.** They point at hosts decommissioned years ago and exist only as attack surface.
4. **Where a static password must remain, set 25 characters or more, randomly generated**, and record a documented exception with an owner and a review date.
5. **Force AES-only ticket encryption on every service account.** Eleven of the fourteen were recoverable because RC4 was still permitted for them. This one setting removes the cheap version of the attack estate-wide and is independent of password strength.
6. **Alert on service tickets requested for Tier-0 adjacent accounts**, separately from the volume-based rule that already exists.

Trade-off: converting to managed accounts requires the service, task, SPN, connection-string, and owner map in plan item L2a. Asset reconciliation in L4 is an input, not that dependency map.

RETEST

`TC-03` · Every account carrying a service principal name either uses a managed service account or has a password set after the remediation date with a length of 25 characters or more; `svc-bkp-legacy` holds no Tier-0 group membership; and no password is recovered during the bounded 72-hour run defined in the client companion runbook. That runbook must record hardware, software, candidate sets, phase counts, elapsed times, and cost before execution. The public specimen does not claim equivalence to the incompletely recorded April run.

REFERENCES

Microsoft, group Managed Service Accounts · T. Medin, original Kerberoasting research, 2014 · MITRE ATT&CK T1558.003 · NIST SP 800-63B section 5.1.1 (memorized secret requirements)

OG-2026-0143-05 display alias F-05

The disbursement API approves payments without an approver identity

CRITICAL PATH IMPACT | BROKE AT DETECTION | NO PIC

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
API APPROVAL	ABSENT	PARTIAL	SILENT	NOT REACHED	EV-19 · 20

SCOPE

The CLARA claims disbursement platform, integration API, production instance. Objective OBJ-1. Primary behaviour `B-31`. Distribution: 1 broke at detection.

ADVERSARY ACTION

Using the integration service account credential taken from the register in F-06, we created a disbursement instruction against a designated synthetic test policy and moved it to the approved state through the API. No human approved it. The instruction was valued at one euro by prior agreement with Finance and was reversed the same day, with written confirmation.

CONTROL

ABSENT. The user interface enforces segregation of duties correctly. The integration API accepts an approval assertion from the service account and does not require an approver identity, because when the interface was built in 2021 the only caller was a trusted internal system. The control exists; it is simply not on the path an adversary uses.

TELEMETRY

PARTIAL. The API call is logged with a timestamp, the service identity, and the resulting state. It does not record an approver, an initiator, or a channel. A legitimate approval and a fraudulent one produce byte-identical log lines. This is judged a detection break rather than a telemetry break because a rule could have been written against the pattern that exists, an integration identity approving outside its normal volume and hours, and was not.

DETECTION

SILENT. No content covers the claims platform. It is not in the detection engineering backlog because the platform is not in the SIEM's owned-source list.

DECISION

NOT REACHED by the SOC. The first signal of any kind arrived at 21:53 on day 31, four days and six hours after the objective, from an overnight access review batch job that flagged the integration account's API pattern. It was raised as an application ticket, not a security event.

CONSEQUENCE

Anyone holding the integration credential can move money without a second person. Read together with F-06, the population holding that credential is 1,340 accounts. Nothing contained it, and the first signal of any kind came 4 days 6 hours after the money had already moved, which means an adversary completes, banks, and is gone long before anything looks.

ROOT CAUSE

Not a persistent catalogue issue. This is an implicit trust boundary from the 2021 integration build: the API was designed for a caller assumed to be trustworthy, and that assumption was never revisited when the credential's blast radius grew. Recurrence risk is **RC-1**, because the fix is a change to a release train that cannot be accelerated.

OWNER

CLARA product owner (approval logic) · Application Security (design review) · Security Engineering (detection content)

REMEDIATION

1. **Require an approver identity on the approval endpoint**, and validate that it is a real user, that it differs from the initiator, and that it is entitled to approve at that value band.
2. **Log initiator, approver, and channel on every state transition.** Do this first even if item 1 is release-bound: without these fields no detection is possible, and with them a compensating detection can ship immediately.
3. **Alert on approvals where the initiator equals the approver, or the approver is a service identity.** This is the compensating control while item 1 waits for its release.
4. **Alert on disbursement state transitions by an integration identity outside its established pattern**, by volume, hour, and value band.
5. **Move the integration credential to a managed secret with short-lived tokens**, removing the static credential that made F-06 fatal.

RETEST

`TC-05` · A disbursement cannot reach the approved state through the API without a distinct, entitled, human approver identity; *and* every state transition log line contains initiator, approver, and channel; *and* an approval by a service identity generates an alert within 15 minutes.

REFERENCES

OWASP API Security Top 10, API5 Broken Function Level Authorization · CWE-863, Incorrect Authorization · No ATT&CK technique describes this business-logic action cleanly, so no technique is assigned.

OG-2026-0143-11 display alias F-11

Tier-0 adjacent alerts are closed without verification

HIGH PATH IMPACT | BROKE AT DECISION | NO PIC | RC-2 | T1558.003 | T1003.006

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
4 BEHAVIOURS	ABSENT	COMPLETE	FIRED	MISSED	EV-12 · 16 · 21

SCOPE	Alert triage process, security operations centre. Four alerts across the operation fired correctly and were closed without a containing action: bulk service ticket requests (day 17), directory replication from a non-domain-controller source (day 20), lateral authentication with a reused local administrator credential (day 24), and scheduled task persistence on a member server (day 36). Primary behaviour `B-13`; supporting `B-16`, `B-22`, `B-29`. Distribution: 4 broke at decision.
ADVERSARY ACTION	None specific to this finding. The finding is what happened to the alerts our other behaviours generated.
CONTROL	ABSENT. No control governs alert closure. Any analyst may close any alert with a free-text reason and no second review, including for assets in the Tier-0 boundary.
TELEMETRY	COMPLETE. All four behaviours were fully recorded. The data was not the problem.
DETECTION	FIRED. All four rules fired correctly and promptly, three of them within 90 seconds. The detection engineering team built the right content. This is worth stating plainly because the finding could otherwise be read as a criticism of them.
DECISION	MISSED, FOUR TIMES. Closure reasons were *Vulnerability scanner — expected* (4 minutes), *backup software* (42 minutes), *imaging activity*, and *known admin tooling*. In no case was the asserted benign cause verified with the team that would have owned it. The scanner assertion was checkable in under a minute: the scanner runs on Tuesdays and this was a Wednesday.
CONSEQUENCE	Four independent opportunities to interrupt the operation were created by the defensive programme and discarded. This is the only finding where the capability was already funded, deployed, and working. The remediation below is process and review time rather than new tooling.
ROOT CAUSE	**RC-2.** Analysts had no feedback loop telling them what these techniques look like when real, and no mechanism made a dismissal visible to the people who wrote the rule. The SOC also ran the entire window with two of eight seats vacant, which is context, not excuse, and it is recorded in section 05.
OWNER	SOC Manager (triage process) · Detection Engineering (feedback loop and runbooks)

REMEDIATION

1. **Replace free-text closure with structured reasons**, and require a verification note naming the source consulted when a benign cause is asserted. "Scanner" becomes "scanner, confirmed with the scanning schedule, ticket #".
2. **Require a second review for any dismissal of an alert touching the Tier-0 boundary.** Four-eyes on closure, not only on escalation.
3. **Return closure outcomes to detection engineering weekly.** The team that wrote these rules did not know they were being dismissed. Give them a dismissal rate per rule and let them fix the ones that read as noise.
4. **Write a runbook per Tier-0 adjacent rule** stating what to check and who to call before closing. Three of the four dismissals were disprovable in under five minutes by someone who knew where to look.
5. **Sample closures.** Ten percent weekly, reviewed by a senior analyst, with the results discussed rather than scored.
6. **Fill the two vacant seats.** Five of the six remediation items above add review steps to a team already running at six of eight. Process changes that assume capacity that is not there will not hold.

This finding is written about a process, not about people. During close-out the analyst who made the day 17 decision walked through it with us, and the reasoning was sound given what they had. The remediation above is what they said would have changed the outcome.

RETEST

`TC-11` · A re-run of the day 17 and day 20 behaviours produces alerts that are either actioned, or closed with a structured reason naming a verified source, within the shift they fired in; *and* a sampled review of the preceding month's Tier-0 adjacent closures shows a verification note on every one.

REFERENCES

MITRE ATT&CK T1558.003, T1003.006, T1078.003, T1053.005 · CIS Critical Security Control 17 (Incident Response Management)

OG-2026-0143-02 display alias F-02

Name-resolution fallback protocols are enabled on a segment nothing logs

HIGH PATH IMPACT | BROKE AT TELEMETRY | PIC-04 | RC-1 | T1557.001

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
NAME POISONING	ABSENT	ABSENT	NOT REACHED	NOT REACHED	EV-09 · 10

SCOPE	`10.90.0.0/16`, the retained branch segment, 412 live hosts, scheduled for decommission since 2024 and directly reachable from the virtual desktop estate. Primary behaviour `B-09`; supporting `B-08`, `B-10`. Distribution: 2 broke at telemetry, 1 not applicable.
ADVERSARY ACTION	We answered name-resolution broadcasts on the segment for eleven hours and collected authentication material from four user accounts and two machine accounts. A mistyped share path from any host there is answered by whichever device replies first, and these protocols do not verify who is replying.
CONTROL	ABSENT. Multicast and broadcast name resolution remain enabled on the segment. Separately, the segment is routable from the virtual desktop estate, which is what let a claims administrator reach it at all.
TELEMETRY	ABSENT. Hosts on the segment do not forward event logs. Local logs had rolled by the time we reconciled records on 15 April, which is why two timeline entries in section 07 are approximate. The wider consequence is that an incident on this segment could not be reconstructed after the fact.
DETECTION	NOT REACHED.
DECISION	NOT REACHED.
CONSEQUENCE	Any host on the segment can be induced to authenticate to an arbitrary listener, and the resulting material can be relayed or recovered offline. SMB relay was blocked by message signing (H-06). The certificate web enrollment attempt also failed (H-10), but the retained IIS configuration and rejection evidence do not establish which control caused it. The collected material did not produce further access on the two tested surfaces. This finding remains High because the exposed segment produced reusable authentication material and no queryable record, while the failed HTTP attempt is only a bounded negative result.
ROOT CAUSE	**RC-1.** Decommission has been deferred twice since 2024 because branch cutover requires downtime nobody will schedule. The protocols themselves are **RC-4**: they are on by default and nothing turned them off.
OWNER	Network Engineering (routing and decommission) · Infrastructure (policy) · Security Engineering (log forwarding)

REMEDIATION

1. **Disable multicast name resolution and NetBIOS name resolution by policy, estate-wide.** Verify the DNS suffix search order first: applications that depend on short-name resolution will break, and finding them afterwards is worse than finding them first.
2. **Remove routing from the virtual desktop estate to this segment.** A user desktop should not reach a decommission-pending branch network. This alone removes the path used in this engagement.
3. **Forward event logs from the segment, or document that it cannot be investigated** and accept that in writing. One of those two things must be true and currently neither is.
4. **Set a dated decommission commitment.** A third year of mitigation costs more than the cutover.

RETEST

`TC-02` · From a host on `10.90.0.0/16`, a request for a non-existent name produces a response only from the authoritative DNS server, and a listener run for four hours collects no authentication material; *and* the segment is unreachable from the virtual desktop estate.

REFERENCES

Microsoft guidance on disabling LLMNR and NetBIOS over TCP/IP · MITRE ATT&CK T1557.001 · CIS Benchmark for Windows, name resolution settings · *A Decade of Security Assessments*, SANS 2025, "Protocols that Refuse to Die"

OG-2026-0143-07 display alias F-07

A monitoring appliance still uses vendor default credentials

MODERATE PATH IMPACT | BROKE AT TELEMETRY | PIC-07 | PIC-12 | RC-5 | T1078.001

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
DEFAULT LOGIN	ABSENT	ABSENT	NOT REACHED	NOT REACHED	EV-17

SCOPE

One network monitoring appliance on the datacentre segment, administrative web interface, reachable from the corporate network without restriction. Primary behaviour `B-25`. Distribution: 1 broke at telemetry.

ADVERSARY ACTION

We authenticated to the administrative interface with the credentials printed in the vendor's public installation guide. The interface then provided a live topology of the estate that is more accurate and more current than the asset inventory the client supplied us at scoping.

CONTROL

ABSENT. Default credentials unchanged since installation in 2023. No network restriction on who can reach the management interface. No second factor.

TELEMETRY

ABSENT. The appliance is not forwarding logs. The device maps the estate but sits outside the monitoring coverage it helps provide.

DETECTION

NOT REACHED.

DECISION

NOT REACHED.

CONSEQUENCE

Reconnaissance quality, mainly. The appliance does not hold credentials to other systems, so the direct impact is bounded. Rated Moderate on what was demonstrated: topology disclosure, reachable from any corporate host, and nothing further. Appendix B says impact is rated on the condition proved rather than a theoretical worst case, and that rule applies here.

ROOT CAUSE	**RC-5.** The appliance is not in the asset inventory (F-12), so it was never in a hardening scope, a patching scope, or a credential review.
OWNER	Network Engineering (device) · Infrastructure (inventory) · Security Engineering (log onboarding)
REMEDIATION	1. **Change the credentials and bind administration to a named account**, federated to the identity provider if the appliance supports it. 2. **Restrict management interface access to an administrative network.** No corporate desktop should reach it. 3. **Onboard the appliance's logs** and alert on administrative authentication. 4. **Sweep for the same pattern.** The verified default on this appliance justifies checking all network and facilities devices. The sweep is scoped in section 11 as a single task.
RETEST	`TC-07` · The vendor default credentials are rejected; the management interface is unreachable from a corporate desktop; an administrative authentication to the appliance appears in the SIEM within 15 minutes; and the estate sweep reports zero devices accepting documented defaults.
REFERENCES	MITRE ATT&CK T1078.001 · CWE-1392, Use of Default Credentials · CIS Critical Security Control 4 (Secure Configuration)

OG-2026-0143-08 display alias F-08

Local administrator credentials are reused across 43 legacy hosts

HIGH PATH IMPACT | BROKE AT DECISION | PIC-10 | RC-5 | T1078.003

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
LATERAL AUTH	DEGRADED	COMPLETE	FIRED	MISSED	EV-16

SCOPE

43 workstations in a legacy imaging group, out of 1,247 workstations in total. Managed local administrator passwords cover the remaining 1,204 (96.6%). Primary behaviour `B-16`. Distribution: 1 broke at decision.

ADVERSARY ACTION

We recovered one local administrator credential and authenticated with it to 42 further hosts. The credential was set by an imaging process that predates the managed password rollout and was never brought into it.

CONTROL

`DEGRADED.` This is a coverage gap in a control that works, not an absent control. The distinction matters for how it is fixed and for how it is reported to a board: the design was right and the rollout has a tail.

TELEMETRY

`COMPLETE.` Every authentication was recorded with source and destination.

DETECTION

`FIRED.` A rule for local administrator authentication from an unexpected source fired on day 24.

DECISION

`MISSED.` Closed as *imaging activity* without checking with the desktop team, who were not imaging that week. See F-11.

CONSEQUENCE

One host compromise becomes 43. In this engagement the group was not a route to anything further, because none of the 43 held privileged sessions during the window. That is luck rather than architecture, and the same reuse in a different group would be a different report.

ROOT CAUSE

RC-5. The imaging group is not tracked as a distinct population, so the managed password rollout reported 96.6% success and nobody asked which hosts made up the remainder.

OWNER

End User Computing (imaging and rollout) · SOC (the triage decision, tracked under F-11)

REMEDIATION

1. **Bring the 43 hosts into managed local administrator passwords** and fix the imaging process that excluded them, so the next batch is not created the same way.
2. **Report rollout coverage as a named exception list, not a percentage.** 96.6% hid a usable attack path; a list of 43 hostnames would not have.
3. **Deny network authentication for local accounts** where the estate can support it, which removes the technique rather than the credential.
4. **Attach a runbook to the existing detection** instructing the analyst to confirm imaging activity with End User Computing before closure.

RETEST	`TC-08` · Every workstation, with no exceptions, holds a unique managed local administrator password; authentication with a previously recovered credential fails on all 43 hosts; and a re-run generates an alert that is actioned or closed with a verified reason within the shift.
REFERENCES	Microsoft, Local Administrator Password Solution · MITRE ATT&CK T1078.003 · CIS Critical Security Control 5 (Account Management)

OG-2026-0143-09 display alias F-09

The password policy reliably produces predictable secrets

MODERATE PATH IMPACT | HELD AT CONTROL | PIC-01 | RC-3 | T1110.003

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
SPRAY + OFFLINE	HELD	COMPLETE	LATE (41 MIN)	NOT REACHED	EV-08 · 11

SCOPE	Domain password policy: 12 characters, complexity enforced, 90-day rotation. 1,180 accounts sprayed; 11 of 14 service account passwords recovered offline. Primary behaviour `B-04`; supporting `B-10`, `B-14`. Distribution: 1 held at control, 2 not applicable.
ADVERSARY ACTION	A low-and-slow spray using season-and-year and company-name-and-year patterns matched three accounts. Offline recovery of captured and requested material broke 11 of 14 service account passwords. The recovered passwords follow the pattern the policy produces: a capitalized word, digits, and a terminating special character, incremented at each rotation.
CONTROL	`HELD, NARROWLY.` The second factor stopped every sprayed account, and that is why this is Moderate rather than High. It held for interactive accounts. It does not apply to service accounts, which is exactly where the policy's weakness became F-03.
TELEMETRY	`COMPLETE` for primary behaviour B-04. The two offline recovery behaviors have no defensive surface and are not applicable; they support the impact finding but do not create a telemetry failure.
DETECTION	`LATE.` The spray detection fired 41 minutes after the spray completed. Not material here because the control held, but the same latency against a technique the control does not cover would be.
DECISION	`NOT REACHED.` The alert arrived after the activity ended and was closed as historical.
CONSEQUENCE	Where a second factor exists, the tested spray was contained. The engagement also recovered 11 service-account passwords offline, but the incomplete compute record does not support a general recovery-time claim. The recovered set followed a rotation-derived increment pattern, which is evidence to remove scheduled rotation and screen new passwords, not a claim about every account.
ROOT CAUSE	**RC-3.** People cannot hold many long complex secrets and rotate them on a schedule, so they produce a formula instead. The policy is doing what it was designed to do and the design is the problem.
OWNER	Identity & Access

REMEDIATION	1. **Remove scheduled rotation for user accounts** and rotate on evidence of compromise instead, in line with current NIST guidance. This removes the increment pattern. 2. **Screen new passwords against a breached-password corpus and a company-specific deny list** containing the organization name, product names, and local seasonal terms. 3. **Raise the minimum length and drop composition rules.** Length beats complexity and does not push people toward formulas. 4. **Extend phishing-resistant authentication** so the control that held here covers more of the estate. 5. **Reduce spray detection latency to under five minutes.** The rule works; its window is too wide.
RETEST	`TC-09` · A newly set password matching the observed pattern is rejected at the point of setting; scheduled rotation is disabled for user accounts; and a repeat spray produces an alert within five minutes.
REFERENCES	NIST SP 800-63B section 5.1.1.2 (memorized secret verifiers; no composition rules, no arbitrary rotation) · MITRE ATT&CK T1110.003 · *A Decade of Security Assessments*, SANS 2025, on rotation making outcomes worse

OG-2026-0143-10 display alias F-10

IPv6 is unmanaged on the legacy segment

MODERATE PATH IMPACT | BROKE AT TELEMETRY | PIC-05 | RC-4 | T1557.003

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
ROGUE DHCPV6	ABSENT (LEGACY)	ABSENT	NOT REACHED	NOT REACHED	EV-24

SCOPE

`10.90.0.0/16` only. The modernized core blocked the identical technique and is recorded as held (H-09). Primary behaviour `B-27`; supporting `B-26`. Distribution: 1 broke at telemetry, 1 held at control.

ADVERSARY ACTION

Two distinct actions, both unfiltered on this segment. We sent DHCPv6 ADVERTISE and REPLY messages offering an operator-controlled DNS server, which hosts accepted, and separately sent unsolicited router advertisements to install an operator-controlled default route. Hosts preferred the IPv6 path and directed name resolution to us, producing forced authentication material. On the modernized core the same two actions were dropped by DHCPv6 Guard and RA Guard respectively before reaching any host.

CONTROL

ABSENT ON THE LEGACY SEGMENT. Present and effective everywhere else, which is what makes this a scoping gap rather than a design failure.

TELEMETRY

ABSENT. Same cause as F-02: the segment forwards nothing.

DETECTION

NOT REACHED.

DECISION

NOT REACHED.

CONSEQUENCE

A second route to the same outcome as F-02 on the same segment. It is rated Moderate rather than High because F-02 already provides that outcome more easily, so this changes the number of ways in rather than what is reachable. It matters because closing F-02 without closing this leaves the path open.

ROOT CAUSE

RC-4. IPv6 is enabled by default on every modern interface, including on a network that has never used it. The switches on this segment predate the guard features deployed on the core.

OWNER

Network Engineering

REMEDIATION

1. **Where the hardware supports it, enable both DHCPv6 Guard and RA Guard on the legacy segment.** They are separate features covering separate message types and one without the other leaves half the path open. Two of the seven switches support both; five support neither.
2. **Where it does not, disable IPv6 on the hosts by policy**, having confirmed nothing on the segment uses it. Nothing does.
3. **Fold this into the decommission commitment in F-02** rather than treating it as separate work.

RETEST	`TC-10` · Neither a rogue DHCPv6 REPLY offering a DNS server nor an unsolicited router advertisement results in any host on `10.90.0.0/16` adopting the offered configuration, verified from three hosts across different switches.
REFERENCES	MITRE ATT&CK T1557.003 (DHCP spoofing) · RFC 6105 for RA Guard; vendor DHCPv6 Guard documentation for the DHCPv6 half · *A Decade of Security Assessments*, SANS 2025, on IPv6 enabled by default

---

OG-2026-0143-14 display alias F-14

Users can grant an application ongoing mail and file access without review

HIGH PATH IMPACT | BROKE AT DETECTION | NO PIC | RC-5 | T1671 | T1550.001

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
CONSENT GRANT	ABSENT	COMPLETE	SILENT	NOT REACHED	EV-29

SCOPE	The Entra ID tenant. Primary behaviour `B-33`. Distribution: 1 broke at detection. User consent to applications is permitted tenant-wide with no admin consent workflow and no publisher verification requirement.
ADVERSARY ACTION	Holding the claims administrator session from day 14, we registered a public-client application, used the authorization-code flow with PKCE, granted delegated mail and file permissions on that user's behalf, and obtained an access token and a user-delegated refresh token. The grant was made through the compromised user session, not by a separate victim or administrator. EV-29 records the client type, flow, delegated permissions, and token class. No password reset or session-revocation test was performed during the operation.
CONTROL	`ABSENT.` User consent is unrestricted. There is no admin consent workflow, no restriction to verified publishers, and no limit on which delegated permissions a user may grant.
TELEMETRY	`COMPLETE.` The consent grant and the service principal creation are both in the tenant audit log with the actor, the application, and the permissions requested. Nothing is missing.
DETECTION	`SILENT.` No content covers consent grants. The data was there, in the platform the detection team already uses, and nothing was watching it.
DECISION	`NOT REACHED.`
CONSEQUENCE	The consent grant remains until it is revoked, while access-token and refresh-token validity is a separate question that depends on token class, client type, password-reset path, and explicit session revocation. Re-registering an authentication factor does not itself remove the grant. This engagement did not test password-reset survival and makes no categorical claim about it. The finding is High because the performed action produced ongoing delegated mail and file access for one user, not tenant-wide application permissions.
ROOT CAUSE	**RC-5.** Nobody owns the review of tenant consent grants, so the question of whether the default should stand was never asked.

OWNER	Identity & Access (tenant policy) · Detection Engineering (content)
REMEDIATION	1. **Restrict user consent** to applications from verified publishers requesting low-impact permissions, and route everything else through an admin consent workflow. 2. **Alert on consent grants and service principal creation.** The telemetry already exists; this is detection content, not a logging project. 3. **Add grant review and explicit session revocation to the account compromise runbook.** Password and factor remediation alone does not establish that delegated application access has ended. 4. **Review existing grants** and remove those with no owner. Expect more than you think.
RETEST	`TC-14` · A standard user cannot grant an application access to mail or files without admin approval; a consent attempt produces an alert within 15 minutes; and a controlled public-client authorization-code/PKCE grant can refresh before, but not after, the runbook's explicit session-revocation step. The record names the client type, flow, token class, revocation mechanism, and both token-use results.
REFERENCES	MITRE ATT&CK T1671 (Cloud Application Integration), with T1550.001 for use of the resulting application access token · Microsoft guidance on user consent settings, refresh-token revocation, and admin consent workflow · CIS Critical Security Control 6 (Access Control Management)

`OG-2026-0143-13` display alias F-13

Inbox rules can hide approval mail, and nothing watches them

MODERATE PATH IMPACT | BROKE AT DETECTION | NO PIC | RC-5

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
RULE CREATION	ABSENT	COMPLETE	SILENT	NOT REACHED	EV-29

SCOPE	The Entra ID tenant, all mailboxes. Primary behaviour `B-32`. Distribution: 1 broke at detection.
ADVERSARY ACTION	From the claims administrator session we created an inbox rule that moved messages matching finance keywords to a folder in the same mailbox and marked them read, so they no longer appeared as unread in the inbox. No message was forwarded or redirected to another recipient. The rule was removed at close-out.
CONTROL	`ABSENT FOR THE TESTED ACTION.` The transport-level block on external auto-forwarding did not prevent or constrain a same-mailbox move and read-mark rule, so this finding gives it no control credit.
TELEMETRY	`COMPLETE.` Rule creation is recorded in the unified audit log with the actor and the rule definition.
DETECTION	`SILENT.` No content covered a rule that moves matching messages and marks them read, although the rule definition was already collected.
DECISION	`NOT REACHED.`

CONSEQUENCE	An adversary already controlling the mailbox can make matching approval messages less visible to that account holder by moving them and marking them read. The test did not forward mail, execute code, or prove access to another mailbox. The finding remains Moderate because approval notifications are part of the claims workflow, but its demonstrated effect is concealment inside one compromised mailbox.
ROOT CAUSE	**RC-5.** The same absent ownership as F-14. Tenant configuration is nobody's review.
OWNER	Identity & Access · Detection Engineering
REMEDIATION	1. **Alert on inbox rule creation** where the rule deletes, marks read, or moves to an obscure folder. This is well-trodden detection content and the telemetry is already there. 2. **Report on rules matching finance and approval keywords** across the tenant as a one-off sweep, then on a schedule. 3. **Keep the external forwarding block as a separate email control**, but do not count it as remediation for the same-mailbox concealment tested here.
RETEST	`TC-13` · Creating a rule that marks matching mail read and moves it to another folder in the same mailbox produces an alert within 15 minutes with the actor, mailbox, and rule actions.
REFERENCES	CIS Critical Security Control 9 (Email and Web Browser Protections). No ATT&CK technique is assigned: T1137.005 requires an Outlook rule used for code execution, and T1114.003 requires forwarding to another recipient; neither was performed.

OG-2026-0143-12 display alias F-12

The asset inventory does not describe the estate

LOW PATH IMPACT | BROKE AT DETECTION | PIC-09 | RC-5 | T1087.002

ACTION	CONTROL	TELEMETRY	DETECTION	DECISION	EVIDENCE
ENUMERATION	ABSENT	PARTIAL	SILENT	NOT REACHED	EV-25

SCOPE	The configuration management database as supplied at scoping, compared against what we enumerated. Primary behaviour `B-12`. Distribution: 1 broke at detection.
ADVERSARY ACTION	We enumerated the directory and the network and compared the result to the inventory provided. 218 live hosts were absent from it, including the appliance in F-07. 96 inventory entries pointed at systems that no longer exist, including three of the service principal names in F-03. Two applications holding regulated data had no recorded owner. The same absent ownership is the root cause behind F-13 and F-14, which are reported separately because they are substantive findings in their own right rather than illustrations of this one.
CONTROL	ABSENT. No reconciliation process compares the inventory to discovered reality.
TELEMETRY	PARTIAL. Directory query content is not logged, so enumeration of this kind cannot be detected. Network discovery is visible.

DETECTION	SILENT.
DECISION	NOT REACHED.
CONSEQUENCE	Rated Low on direct path impact because no single step of the compromise depended on it. It is placed last in the register and near the top of the improvement plan, which looks contradictory and is not: on its own it grants an adversary nothing, and it is the reason four other findings existed. Hardening, patching, and monitoring scopes are all derived from this inventory, so anything missing from it is outside all three by construction.
ROOT CAUSE	**RC-5**, definitionally. This is the root cause the catalogue names, seen directly rather than through one of its consequences.
OWNER	Infrastructure (inventory) · Enterprise Architecture (ownership records)
REMEDIATION	1. **Reconcile the inventory against authenticated discovery on a schedule**, and treat the difference as a work queue rather than a report. 2. **Require a named owner for every application holding regulated data**, with no unowned entries permitted to persist past a review cycle. 3. **Feed the discovery output into the hardening, patching, and log-onboarding scopes automatically**, so a newly discovered asset joins all three without a human deciding to add it. 4. **Enable directory query logging on domain controllers** so enumeration at this scale becomes detectable, accepting the volume cost. 5. **Review tenant application consent grants and mailbox rules** and assign an owner to both.
RETEST	`TC-12` · A discovery run reconciles to the inventory with under 2% unexplained difference; every application holding regulated data has a named owner; and a newly introduced test host appears in the inventory and in the log-onboarding scope within one reconciliation cycle.
REFERENCES	CIS Critical Security Controls 1 and 2 (Inventory of Enterprise Assets and Software) · MITRE ATT&CK T1087.002

10 RESULTS

10 What held

Seven controls and two response processes stopped a technique we ran deliberately to test them, and one entry (H-01) is a bounded negative result rather than a control hold. All ten are recorded here. These are recorded to the same evidence standard as the failures, for a specific reason: a report that lists only breaks gives leadership no way to distinguish a control that works from one that was never tested. Both look identical on a page that does not mention them.

Everything below is a statement about a specific behaviour at a specific time, not a general assurance. "Held" means the recorded behaviour was stopped, and the behaviour is named so the claim can be checked and re-run.

CONTROLS THAT STOPPED A TESTED BEHAVIOUR

ID	WHAT HELD	BEHAVIOUR	EVIDENCE AND LIMIT OF THE CLAIM
H-01	**No exploitable path was found on the internet-facing estate** A bounded negative result, not a control hold	B-01 (contextual result)	214 responsive services enumerated across both declared ranges over eight days. No unauthenticated request changed state on any of them, and no exposed management interface was unpatched against a publicly known issue. *Read this precisely.* It says we did not find a way in during that window. It is not a control stopping an attack, and section 05 is explicit that this method does not prove absence. It is recorded because a buyer should be able to tell a tested-and-clear surface from an untested one, and because it is why initial access had to go through a person.
H-02	**The print spooler is disabled on domain controllers**	B-21	Coerced machine authentication against all four domain controllers produced no connection. The service is disabled by policy on the Tier-0 build. *Limit:* it remains enabled on member servers, where the same coercion succeeds; that is lower impact and is in the plan rather than the findings.
H-03	**Database tier segmentation defeated coerced authentication**	B-28	Coercion from a domain-joined database server succeeded in producing authentication material, and the material could not reach operator infrastructure because those hosts have no route out. The technique worked; the architecture made it worthless. *Limit:* a listener placed inside the datacentre segment would collect it, so this holds against an external position only.
H-04	**Conditional access and the second factor defeated stolen passwords**	B-03, B-04	Eleven phished credentials and three sprayed matches produced zero sessions. Every attempt was refused with the correct policy reason recorded. *Limit:* this held against possession of a password. It did not hold against control of the factor itself, which is F-01, and that is the distinction the organization should take from this pairing.
H-05	**Data egress was detected and contained in eleven minutes**	B-30	Two independent detections fired at +5 m 55 s, an analyst correlated them and terminated the session at +11 m 06 s on their own authority, and the account was disabled at +23 m 34 s. No records left the platform. The full reconciled timeline is EV-22. This is the only objective-path behaviour actioned before its objective. B-02 also ran through every chain link, but no engagement objective depended on it.

ID	WHAT HELD	BEHAVIOUR	EVIDENCE AND LIMIT OF THE CLAIM
H-06	**SMB message signing is enforced across the estate**	B-11	Relay attempted against 61 hosts holding collected authentication material. All 61 refused. *Limit:* signing protects SMB and nothing else. It does not make captured material unusable against HTTP or LDAP endpoints. The separate HTTP enrollment attempt in H-10 failed, but the retained evidence does not establish which IIS control caused that rejection.
H-07	**Kerberos pre-authentication is required on every account**	B-15	A full sweep of the directory found no account with pre-authentication disabled, so the offline attack that requires it had no targets. The result is bounded to the directory snapshot collected for that sweep.
H-08	**Staff reported phishing, and email response acted on it**	B-02	19 reports from 218 delivered messages, the first at four minutes. The sending domain was blocked at two hours six minutes. The reporting rate exceeded the credential submission rate in this exercise.
H-09	**DHCPv6 Guard held on the modernized core**	B-26	We sent DHCPv6 ADVERTISE and REPLY messages offering an operator-controlled DNS server. DHCPv6 Guard dropped them at the access switch before any host saw them. RA Guard is configured on the same switches and unsolicited router advertisements were dropped as well. The identical behaviour succeeded on the legacy segment (F-10), which makes this a clean demonstration that the modernization programme produces the intended result where it has reached.
H-10	**The enrollment endpoint rejected the tested HTTP relay**	B-36	The certificate authority's web enrollment endpoint is reachable over HTTP with NTLM, which is the textbook target for relaying captured authentication material into a certificate request. We tried it. The relayed request was rejected and no certificate was issued. EV-31 did not preserve the IIS `tokenChecking` value, flags, SPNs, Windows Authentication provider order, request template, or HTTP/IIS status. The report therefore cannot attribute the rejection specifically to Extended Protection; a request, tool, template, or service-binding mismatch remains possible. *Limit:* this is only a bounded failed attempt. Require HTTPS, disable NTLM, record the complete IIS configuration, and retest the existing and hardened states before crediting a named control.

THE PATTERN IN WHAT HELD

Eight matrix rows held at control: four identity and directory rows (`B-03` , `B-04` , `B-15` , `B-36`), three network rows (`B-11` , `B-26` , `B-28`), and one endpoint row (`B-21`). Two more rows were actioned: email behavior `B-02` and application behavior `B-30` . H-01 is a bounded negative result, not a hold. B-30 is the only objective-path behavior actioned before its objective.

10.4 Persistent issue coverage

Every issue in the catalogue, and what this engagement found. Entries marked *tested, not present* are as much a result as the findings: they say we looked.

ID	PERSISTENT ISSUE	COVERAGE	RESULT
PIC-01	Predictable passwords from complexity-and-rotation policy	TESTED: PRESENT	Pattern confirmed across sprayed and recovered credentials. F-09.

ID	PERSISTENT ISSUE	COVERAGE	RESULT
PIC-02	Weak MFA methods and weak recovery for strong ones	TESTED: PRESENT	Methods themselves are strong. Recovery is a phone call. F-01.
PIC-03	Neglected internal perimeter behind a maintained external one	TESTED: PRESENT	Structural rather than a single finding. The external estate produced a bounded negative result (H-01); all fourteen findings are internal.
PIC-04	Name-resolution fallback protocols	TESTED: PRESENT	Legacy segment only. Disabled elsewhere. F-02.
PIC-05	Unmanaged IPv6	TESTED: PRESENT	Predicate: unmanaged IPv6 anywhere in the tested network sample. Present on legacy; blocked on the modernized core (H-09). F-10.
PIC-06	Print spooler and coerced machine authentication	TESTED: NOT PRESENT	Predicate: coerced machine authentication exposure on the tested Tier-0 domain controllers. The spooler was disabled on all four (H-02). Member-server exposure was observed outside this bounded predicate and remains a plan item.
PIC-07	Internal portals with no or default authentication	TESTED: PRESENT	One appliance. 34 further internal web interfaces were checked and required authentication. F-07.
PIC-08	Certificate services misconfiguration	TESTED: PRESENT	One template exploited; three others carry the risky purpose with narrower rights. Web enrollment reachable over HTTP with NTLM. F-04.
PIC-09	Sensitive data on shares; inventory failure	TESTED: PRESENT	Both halves. Credential register (F-06) and inventory divergence (F-12).
PIC-10	Password reuse, especially local administrator	TESTED: PRESENT	Predicate observed on 43 hosts; 96.6% of workstations were covered by managed passwords. F-08.
PIC-11	Weak or never-rotated service account passwords	TESTED: PRESENT	11 of 14 recovered, oldest set in 2019, one holding Tier-0 rights. F-03.
PIC-12	Default credentials on appliances and management software	TESTED: PRESENT	One confirmed. A full estate sweep was not in scope and is recommended in section 11. F-07.
PIC-13	SQL Server `sa` and domain-authenticated SQL abuse	NOT IN SCOPE	SQL authorization abuse and `sa` testing were not in the validation plan. B-28 tested network egress after coerced authentication from a database server; it does not establish this PIC predicate.
PIC-14	Active Directory misconfiguration chains	TESTED: PRESENT	The compromise path is one: an enrollment right on a permissive template, reached from a service account with an old password. F-03 and F-04 together.

Twelve predicates were tested and present, one was tested and not present, and one was not in scope. Several present rows rest on bounded samples and do not claim estate-wide prevalence. Seven issues were load bearing in the compromise path: PIC-02 (the way in), PIC-01 and PIC-11 (the service account passwords), PIC-14 and PIC-08 (the certificate chain), PIC-09 (the credential workbook), and PIC-03 as the structural condition underneath them. PIC-04 was present but its captured material unlocked nothing, so it is not counted as load bearing. This is what the phrase "issues that refuse to die" means in practice: nothing in this engagement required a technique that was not publicly documented before 2022, and the oldest of them has been documented for over a decade.

11 FOR LEADERSHIP

11 Improvement plan

Sequenced rather than sorted. Each item names the team that agreed to own it during the ownership session on 11 May, an effort estimate produced by that team rather than by us, and the retest condition that will count as done. Items are ordered so that dependencies resolve before the work that needs them.

EFFORT S < 1 WEEK · M 1–4 WEEKS · L > 4 WEEKS OR RELEASE-BOUND OWNER AGREED AT THE 11 MAY SESSION

Now: within 30 days

Everything here closes a demonstrated path to forest control or to money, and none of it is release-bound.

#	ACTION	CLOSES	OWNER	EFFORT	RETEST
N1	Delete or reconfigure the `NWM-Legacy-WebServer` template; require approval on the three other templates carrying a client authentication purpose.	F-04	Identity & Access	S	TC-04
N2	Remove `svc-bkp-legacy` from the Tier-0 group; grant only the rights the backup job needs.	F-03	Identity & Access	S	TC-03
N3	Remove the credential workbook; rotate all 47 credentials; treat as a credential compromise event, not a finding.	F-06	Identity & Access + Finance	S	TC-06
N4	Change service desk factor-reset procedure to verify through a channel the caller did not choose; prohibit reset entirely for privileged accounts.	F-01	IT Service Management	S	TC-01
N5	Remove the verification-question procedure from public content.	F-01	Corporate Comms	S	TC-01
N6	Replace free-text alert closure with structured reasons; require four-eyes on any dismissal touching the Tier-0 boundary.	F-11	SOC	S	TC-11
N7	Change the monitoring appliance credentials and restrict its management interface to an administrative network.	F-07	Network Engineering	S	TC-07
N8	Remove routing from the virtual desktop estate to `10.90.0.0/16`.	F-02 F-10	Network Engineering	S	TC-02
N9	Log initiator, approver, and channel on every disbursement state transition, ahead of the enforcement change.	F-05	CLARA product	M	TC-05

Next: 30 to 90 days

Telemetry and detection work. This is where strategic move 01 from the executive summary actually happens, and it is the block that closes the eight collectable blind spots while adding detection to source records that already exist.

#	ACTION	CLOSES	OWNER	EFFORT	RETEST
X1	Enable CA security auditing, forward it to the SIEM, monitor the retained CA issuance records, and verify supported strong certificate mapping on each domain controller.	F-04	Identity & Access + Sec Eng	M	TC-04
X2	Forward identity provider events for security-information deletion, Temporary Access Pass creation, and new-method registration; correlate them with an unfamiliar-device sign-in.	F-01	Security Engineering	M	TC-01
X3	Disable multicast and NetBIOS name resolution estate-wide, after verifying DNS suffix search order.	F-02	Infrastructure	M	TC-02
X4	Enable object access auditing on shares holding credential or regulated data; alert on bulk read.	F-06	Infrastructure + Sec Eng	M	TC-06
X5	Bring the 43 legacy imaging hosts into managed local administrator passwords; fix the imaging process; report coverage as an exception list.	F-08	End User Computing	M	TC-08
X6	Deploy breached-password screening with a company-specific deny list; raise minimum length; remove scheduled rotation for user accounts.	F-09	Identity & Access	M	TC-09
X7	Return alert closure outcomes to detection engineering weekly; write runbooks for every Tier-0 adjacent rule.	F-11	SOC + Detection Eng	M	TC-11
X8	Sweep all network and facilities devices for documented default credentials; onboard their logs.	F-07	Network Engineering	M	TC-07
X9	Forward event logs from `10.90.0.0/16`, or record a written acceptance that the segment cannot be investigated.	F-02 F-10	Security Engineering	M	TC-02
X9b	Restrict user consent to verified publishers and low-impact permissions; route the rest through admin consent. Add token revocation to the compromise runbook.	F-14	Identity & Access	S	TC-14
X9c	Build detection content for consent grants, service principal creation, and inbox rule creation. The telemetry is already collected.	F-13 F-14	Detection Engineering	S	TC-13 TC-14

#	ACTION	CLOSES	OWNER	EFFORT	RETEST
X10	Alert on approvals where initiator equals approver or the approver is a service identity, as the compensating control while N9 and L3 land.	F-05	Detection Engineering	S	TC-05

Later: beyond 90 days

Programme work and release-bound changes. Two of these are the reason several findings existed in the first place, which is why they are listed even though nothing here is quick.

#	ACTION	CLOSES	OWNER	EFFORT	RETEST
L1	Provide a secrets manager with delegated access and migrate the register into it. Without this, F-06 returns under a different filename.	F-06 F-05	Identity & Access	L	TC-06
L2a	Build a service-account dependency map: logon and service and scheduled-task inventory per account, SPN ownership, connection strings, application owner validation, and a tested rollback and rotation window.	F-03	Identity & Access + app owners	L	TC-03
L2b	Convert service accounts to managed service accounts once L2a covers them; document exceptions with owners and review dates. Set `PrincipalsAllowedToRetrieveManagedPassword` explicitly per account.	F-03	Identity & Access + app owners	L	TC-03
L3	Require a distinct, entitled, human approver identity on the disbursement API. Release-bound.	F-05	CLARA product	L	TC-05
L4	Reconcile the asset inventory against authenticated discovery on a schedule; feed the output into hardening, patching, and log-onboarding scopes automatically.	F-12	Infrastructure	L	TC-12
L5	Complete the `10.90.0.0/16` decommission against a dated commitment.	F-02 F-10	Network Engineering	L	TC-02 TC-10
L6	Extend phishing-resistant authentication across the remaining estate.	F-09 F-01	Identity & Access	L	TC-09
L7	Fill the two vacant SOC analyst seats.	F-11	Head of Cyber Defence	L	TC-11

THE DEPENDENCY WORTH NAMING

L2a blocks L2b, and L2b is the real fix for a Critical. Service accounts cannot be safely converted to managed accounts without knowing every dependency of each one, and that is a different exercise from the asset reconciliation in L4. L4 tells you which machines exist; it does not tell you which scheduled task on which host authenticates as `svc-print-r2`. Conflating the two would park a Critical behind a programme that never produces its prerequisite, which is why L2a is listed separately. The plan sequences around the risk by doing N2 immediately, removing the Tier-0 membership, so the conversion work can take the time it needs without leaving a Critical open.

What this plan does not include

- **Any purchase.** Every item above is configuration, process, ownership, or engineering time against systems already deployed. If a vendor conversation follows this report, it should be about L1 and nothing else.

- **The member server print spooler.** Enabled outside Tier-0 and exploitable for coercion, but not raised as a finding because message signing (H-06) makes the result unusable. Worth doing; not worth doing before anything above it.
- **The Dublin recovery forest.** Out of scope and untested. It holds a trust with the primary forest, which makes it worth a scoped assessment once the primary forest work is complete. This report does not claim it is affected.

12 FOR LEADERSHIP

12 Retest record

Five days in June, re-running all fourteen documented validation conditions, not only the ones for the Critical and High findings. This section records what was verified, what was verified only in part, what the client chose to accept, and what could not be checked at all.

Retest window	8 – 12 June 2026
Scope	All 14 documented validation conditions, `TC-01` to `TC-14`. Not a re-run of the engagement.
Method	Announced. The SOC knew, which is correct for a retest: the question is whether the fix works, not whether anyone notices.
Evidence	Retest evidence carries the prefix `RV-` and is indexed separately in Appendix E.

What changed

Both demonstrated routes to Tier-0 no longer function. The certificate template is gone and the domain controllers now refuse a weakly mapped certificate (F-04); the Tier-0 service account no longer holds those rights (F-03). The telephone route to an authentication factor is closed (F-01). The four alerts that fired during the operation and were dismissed were actioned within the shift when re-run (F-11). An operator repeating this exact end-to-end operation today would be stopped at initial access, and both tested routes to OBJ-3 would also fail. OBJ-1 remains technically reachable to anyone who already holds the integration credential, but its current detection and cancellation path prevents the payment file from carrying the instruction. Three upstream points now hold or alert where neither happened during the operation.

What remains open is specific rather than general: a payment can still be approved through the API without a human, and is now detected rather than prevented (F-05); the inventory programme has not started (F-12); and one segment carries an accepted risk until November (F-10).

> WHY THERE IS NO SINGLE ENGAGEMENT RATING HERE
>
> A previous version of this report closed with one rating for the whole engagement, moved from Critical to Moderate. It has been removed. There is no defensible rule for aggregating fourteen findings across two independent axes into one number, and inventing one would be the exact failure this report criticises elsewhere. The per-finding table above and the residual position at the end of this section say everything the rating was trying to say, and they say it in a form the client can argue with. Where a risk register requires a single value, it should be derived by the client against their own impact scale, from the finding-level ratings in section 09.

Original finding distribution

COMPONENT	CRITICAL	HIGH	MODERATE	LOW	TOTAL
Full-scope adversary validation	5	4	4	1	**14**

Per-finding retest status

ID	FINDING	STATUS	CHAIN VERDICT NOW	WHAT WAS VERIFIED
F-01	Factor restored by telephone	REMEDIATED	Held at control	An authorized simulated call using only public information was refused and escalated. A control factor change performed with approval appeared in the SIEM in 3 m 40 s and correlated to the subsequent sign-in. RV-01
F-04	Certificate template and unmonitored issuance records	REMEDIATED	Held at control	Template deleted. The three other client-authentication templates now require approval. Issuance auditing enabled and forwarding; a control request produced a SIEM event in 4 m 10 s. Strong certificate mapping enforced. Engagement certificate revoked. RV-04
F-06	Credential register on a broad share	REMEDIATED [1]	Actioned	Workbook absent from all indexed shares. All 47 credentials show a set date after 15 May. A credential-pattern scan returned two low-confidence hits, both reviewed and dismissed with the client. A control bulk read of Finance$ alerted in 7 m. RV-06
F-03	Service account passwords from 2019	PARTIALLY REMEDIATED [2]	Not applicable	svc-bkp-legacy no longer holds Tier-0 rights. 9 of 14 accounts converted to managed accounts. 3 orphaned principal names removed. 2 accounts remain on their original 2019 passwords. RV-03
F-05	Disbursement approved without an approver	PARTIALLY REMEDIATED [3]	Actioned after objective	Initiator, approver, and channel now logged on every state transition. On a control test the detection fired at 6 m 12 s and an analyst cancelled the disbursement instruction at 14 m 03 s. Disbursements are written to the payment file in the 23:00 run, so cancellation before then prevents money moving: the objective state was reached and the harm was not. Enforcement remains release-bound. RV-05

ID	FINDING	STATUS	CHAIN VERDICT NOW	WHAT WAS VERIFIED
F-11	Tier-0 adjacent alerts closed without verification	REMEDIATED	Actioned	Structured closure reasons live; four-eyes required on Tier-0 adjacent dismissals; weekly closure review returning to detection engineering; runbooks written for 11 rules. Re-running the day 17 and day 20 behaviours produced containing actions in 9 and 14 minutes. RV-11
F-14	Unreviewed application consent grant	REMEDIATED	Held at control	User consent restricted to verified publishers requesting low-impact permissions; everything else routes through admin consent. A control grant attempt was refused and alerted in 5 m 20 s. A controlled public-client authorization-code/PKCE grant refreshed successfully before, and failed after, explicit session revocation in the account-compromise runbook. Existing grants reviewed; 11 removed as unowned. RV-14
F-13	Same-mailbox concealment rule unmonitored	REMEDIATED	Actioned	A control rule that marked matching mail read and moved it within the same mailbox alerted in 8 m 45 s and was actioned in 19 m. A tenant-wide sweep found four pre-existing rules matching finance keywords, all reviewed with the account holders. RV-13
F-02	Name resolution on an unlogged segment	REMEDIATED	Held at control	Multicast and NetBIOS name resolution disabled estate-wide. A four-hour listener on the segment collected nothing. Virtual desktop routing to the segment removed. Event forwarding enabled. Decommission committed for 30 November 2026. RV-02
F-07	Appliance with default credentials	REMEDIATED	Held at control	Defaults rejected; administration federated and restricted to an administrative network; logs onboarded and an administrative sign-in alerted. The estate sweep found four further devices with documented defaults, all since remediated. RV-07

ID	FINDING	STATUS	CHAIN VERDICT NOW	WHAT WAS VERIFIED
F-08	Local administrator reuse on 43 hosts	REMEDIATED	Actioned	All 43 hosts hold unique managed passwords; the imaging process is corrected; coverage is now reported as a named exception list. Authentication with the recovered credential failed on all 43, and the resulting alert was actioned in 12 minutes. RV-08
F-09	Password policy produces predictable secrets	PARTIALLY REMEDIATED [4]	Held at control	Breached-password screening and a company deny list are live and reject the observed pattern at the point of setting. Minimum length raised to 15. Scheduled rotation removed for user accounts. Spray detection now fires in under 4 minutes. RV-09
F-10	Unmanaged IPv6 on the legacy segment	RISK ACCEPTED [5]	Broke at telemetry	Not remediated by decision. Accepted in writing by the CISO on 11 June 2026 with a review date of 30 November 2026. RV-10
F-12	Asset inventory does not describe the estate	NOT REMEDIATED [6]	Broke at detection	No change. The reconciliation programme is scoped and funded with a start date of 1 September 2026 and a named owner. RV-12

Notes on the partial and accepted statuses

A status of anything other than *Remediated* is only useful if it says exactly what remains. These notes are the point of the section.

NOTE	FINDING	WHAT IS OUTSTANDING, AND THE RESIDUAL POSITION
1	F-06 · credential register	**The stated retest condition is met.** The recurrence control is not: the secrets manager (plan item L1) covers 31 of 47 applications, so 16 teams still have no supported place to keep a shared credential. We assess recurrence risk as *moderate*, meaning we expect a replacement file to exist within a year unless L1 completes. Recommend a targeted re-scan in the next validation cycle rather than waiting for an annual test.
2	F-03 · service account passwords	Two accounts, `svc-print-r2` and `svc-edi-gw`, remain on their 2019 passwords because their service, scheduled-task, SPN, connection-string, and application-owner dependencies are unknown and the owning teams will not risk an outage. **One of the two was recovered again during the retest** in under 40 hours. Neither holds Tier-0 rights, so the path proven in this engagement is closed, but the underlying issue is live on two accounts. This is blocked on plan item L2a (service-account dependency mapping). L4 supplies asset inventory input but does not produce the required dependency map.
3	F-05 · disbursement approval	A payment can still be approved through the API without a human, and the exact path demonstrated in this engagement remains available to anyone holding the integration credential. What changed is that the approval is now detected and cancelled before the 23:00 payment run. The chain verdict moved from *broke at detection* to *actioned after objective*, and that qualifier is doing real work: the objective is still reached, so this is interception rather than prevention. It is not recorded as a success and F-05 stays Critical until L3 ships on 18 September. The containing action is a manual cancellation inside a nine-hour window; it depends on an analyst being on shift and on the payment run not being brought forward. Enforcement (plan item L3) is scheduled for the release of 18 September 2026. Until then the compensating detection is load bearing, so any change to it must be treated as a change to a Critical control. **We recommend the detection be tested monthly, not annually**, for exactly this reason.
4	F-09 · password policy	New passwords are screened; existing ones are not. Because scheduled rotation was correctly removed at the same time, there is now no event that would cause the existing population to be re-screened, so today's predictable passwords will persist indefinitely. This is a side effect of two individually correct changes. **Recommend a one-time forced reset with screening applied**, then no further scheduled rotation.
5	F-10 · unmanaged IPv6	Five of the seven switches on the segment cannot support the guard features, and disabling IPv6 by policy was tested and broke a branch printing workflow. The client accepts the risk until the segment is decommissioned on 30 November 2026. Compensating controls in place and verified: the segment is no longer reachable from the virtual desktop estate, and its event logs are now forwarded, so the same activity would at least be reconstructable. **Accepted by H. Okonkwo, CISO, 11 June 2026. Review date 30 November 2026.**
6	F-12 · asset inventory	No work has started. The programme is scoped, funded, owned by Infrastructure, and committed to start on 1 September 2026. Recorded as not remediated rather than in progress, because a commitment is not a control. L4 feeds asset data into L2a but does not block F-03 by itself. It remains the root cause named in F-07 and F-08.

What the retest could not verify

Stated because a retest that reports only successes is not a verification.

- **Whether the certificate template was ever abused before this engagement.** The CA database retains issuance records and can be queried, but the retest did not complete a lifetime review and an issued certificate record alone does not establish malicious intent. This remains unverified rather than permanently unanswerable.
- **Whether the credential workbook was read by anyone else.** Same cause: object access auditing was not enabled on the share. All 47 credentials were rotated on that basis, which is the correct response to an unanswerable question.

- **Whether the wider classes of issue are closed.** The retest verified fourteen specific validation conditions. It did not re-run the full 36-behaviour matrix, so the chain verdict statistics in section 08 stand as of April and are not restated here. Ten findings changed their verdict on their own validation condition: F-01, F-02, F-04, F-05, F-06, F-07, F-08, F-11, F-13, and F-14. Those are shown above.
- **Whether the process changes will hold.** F-01 and F-11 are procedural. They were verified in June against a team that knew a retest was running. Procedural controls decay, and the only honest way to know is to test them unannounced in a later cycle.

RESIDUAL POSITION

An operator repeating this exact end-to-end operation today would be stopped at the telephone route to an authentication factor. Both tested routes to Tier-0 also no longer function. Those three upstream points now hold or alert where neither happened during the operation.

That is not the same as every path being closed. **The F-05 payment path remains open** to anyone holding the integration credential, and is intercepted rather than prevented until the September release. Saying otherwise would be the risk laundering this report criticises elsewhere.

What remains is not a set of open holes so much as three commitments that have to be kept: the disbursement enforcement release in September, the inventory programme in September, and the segment decommission in November. Two Critical findings currently depend on a compensating detection or on a future date rather than on a control, and that is the condition to re-examine in the next validation cycle rather than at the next annual test.

Retest performed and this record prepared by Jean-François Maes, 8–12 June 2026. Issued as version 1.0 on 19 June 2026. Findings, evidence, and test cases are retained encrypted until 17 September 2026 and destroyed thereafter unless the client instructs otherwise in writing.

A APPENDIX

A Scope inventory and attribution

Operator infrastructure

Declared before the operation so the client can reconcile independently. All operator activity in the engagement originated from these addresses.

RANGE OR IDENTIFIER	USE	ACTIVE
203.0.113.16/28	External assessment and command infrastructure	2 Mar – 11 Apr 2026
203.0.113.64/29	Phishing campaign delivery and capture	10 Mar – 13 Mar 2026
10.30.44.180	Operator host inside the virtual desktop estate	15 Mar – 11 Apr 2026
+32 · withheld	Voice origination for the day 14 call, spoofed to an internal extension	15 Mar 2026
203.0.113.16/28	Retest, all activity	8 – 12 Jun 2026

In-scope ranges

RANGE	DESCRIPTION	LIVE HOSTS	NOTE
192.0.2.0/24	Primary internet-facing estate	148	214 responsive services in total across both external ranges
198.51.100.0/24	Secondary internet-facing estate	41	Includes the tenant federation endpoints
10.30.0.0/16	Corporate network, user VLANs and virtual desktop estate	1,388	1,247 workstations, 141 infrastructure hosts
10.40.0.0/16	Datacentre, application and database tiers	312	CLARA and POLARIS both reside here
10.90.0.0/16	Legacy branch segment, decommission pending	412	In scope at client request. Source of F-02 and F-10

Address ranges use blocks reserved for documentation (RFC 5737) and private use (RFC 1918). Host and domain names use the reserved `.example` top-level domain (RFC 2606). No real infrastructure is described anywhere in this document.

Inventory divergence

COMPARISON	COUNT	NOTE
Live hosts absent from the supplied inventory	218	Includes the appliance in F-07
Inventory entries with no corresponding live host	96	Includes three service principal names from F-03
Applications holding regulated data with no recorded owner	2	Raised to Enterprise Architecture during the engagement

B APPENDIX

B Rating definitions and CVSS

Path impact

RATING	DEFINITION	EXPECTED RESPONSE
CRITICAL	Achieved or trivially achievable control of the organization's identity, data, or money at scale. Forest compromise, mass data access, control of a payment path.	Same-week action. Treat the associated credentials or identities as compromised.
HIGH	Material unauthorized access or privilege that reaches Critical with routine effort. Tier-0 adjacency, broad credential access, sensitive dataset access.	Action within the current planning cycle. Do not defer past a quarter.
MODERATE	Meaningful unauthorized access, or a required link in a demonstrated chain. Not sufficient alone; load bearing in combination.	Schedule. Re-rate if the finding it depends on is not fixed.
LOW	Genuine weakness with limited standalone consequence in this environment.	Fix in normal maintenance. May still rank high in the plan if it is a root cause.
INFORMATIONAL	Observation with no demonstrated path. Recorded for completeness.	Judgement.

Ratings reflect what was demonstrated in this environment during this window, not a vulnerability's theoretical worst case. A true risk rating cannot be produced by offensive testing alone, because likelihood depends on threat exposure and business context the engagement does not measure. These ratings are an input to risk assessment, not a substitute for it.

Chain verdict

VERDICT	MEANING	WHO FIXES IT
HELD AT CONTROL	A preventive or limiting control stopped or materially changed the path.	Nobody. Protect it and keep testing it.
ACTIONED BEFORE OBJECTIVE	The chain worked and containment landed before the adversary achieved the objective. The defence prevented the outcome. Where no objective was in play for a behaviour, the verdict is written simply as *Actioned*: the qualifier depends on whether there was an objective to be early or late for, not on the node states.	Nobody. This is the target state for everything else.
ACTIONED AFTER OBJECTIVE	The alert fired and someone acted, but the objective had already been met. The defence found out; it did not prevent.	Whoever owns the control that should have stopped it. Detection is working and is not the gap.
NOT APPLICABLE	The behaviour has no defensive surface. Offline password cracking happens on our hardware; reading your public help portal is not an event you can log. Recorded because it happened, excluded from every break count and percentage.	Nobody. Counting these against you would inflate the number and misdirect the spend.

VERDICT	MEANING	WHO FIXES IT
BROKE AT TELEMETRY	No queryable record with enough context existed. No detection was possible.	Platform and log engineering. Collection is the prerequisite for downstream detection.
BROKE AT DETECTION	The activity was recorded and no logic was watching for it.	Detection engineering. The source record exists; the gap is forwarding, analytic content, or both.
BROKE AT DECISION	An alert fired and no containing action followed.	Operations process, runbooks, staffing. No product fixes this.

How a verdict is chosen

A verdict belongs to a behaviour. The verdict runs left to right along the chain and stops at the first node that failed; where two nodes are degraded the earlier one wins, because it had to hold first. A detection that fired *late* still counts as fired, with the latency recorded separately. A detection that has been tuned out counts as silent, because logic that does not fire is logic that does not fire.

A *finding* usually covers several behaviours. It takes the verdict of one designated primary behaviour, named in the finding's scope, and the others are listed alongside. That rule exists so a reader can see which row the headline came from instead of guessing whether it was the worst one, the first one, or the most convenient one.

Percentages in this report are taken over the **assessable** behaviours: the total minus those marked not applicable. Both numbers are always printed.

CVSS v3.1 base scores

Supplied because the client's risk register requires a CVSS value. It is not the headline rating and it is not used to order the plan. Every vector below was recomputed with the CVSS v3.1 formula and the printed score is the formula's output. Note how poorly it separates these findings: two of the four scoreable Criticals (F-05 at 7.1, F-06 at 7.7) score below a High finding (F-08 at 8.0), and F-07, which this report rates Moderate on demonstrated impact, still scores 7.5. CVSS scores a vulnerability in isolation; this engagement scored a path.

ID	PATH IMPACT	CVSS 3.1	VECTOR	NOTE
F-01	CRITICAL	Not scored	Not applicable	Process finding. No software weakness to score.
F-02	HIGH	6.8	AV:A/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:N	
F-03	CRITICAL	8.8	AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H	
F-04	CRITICAL	9.9	AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:H	Scope change reflects forest-wide identity impact.
F-05	CRITICAL	7.1	AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:H/A:N	Re-vectored to `PR:L`. Needing a credential does not make privileges required high, and F-06 put this one in reach of 1,340 accounts.
F-06	CRITICAL	7.7	AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:N/A:N	
F-07	MODERATE	7.5	AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N	Re-vectored: only topology disclosure was demonstrated, so integrity and availability are none. Still scores above two Critical findings.

ID	PATH IMPACT	CVSS 3.1	VECTOR	NOTE
F-08	HIGH	8.0	AV:A/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:H	
F-09	MODERATE	5.3	AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N	
F-10	MODERATE	6.5	AV:A/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N	
F-13	MODERATE	4.3	AV:N/AC:L/PR:L/UI:N/S:U/C:L/I:N/A:N	
F-14	HIGH	9.6	AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:N	No separate victim interaction occurred: the operator acted through an already compromised user session. Scope change reflects delegated access to mail and files through the application boundary.
F-11	HIGH	Not scored	Not applicable	Process finding. CVSS has no way to express "the alert fired and was dismissed".
F-12	LOW	Not scored	Not applicable	Programme finding.

Three of the fourteen findings cannot be scored at all, and they include the finding that granted initial access and the finding that discarded four chances to stop the operation. A report that leads with CVSS would have to either omit them or invent a number for them.

C APPENDIX

C ATT&CK coverage

Every technique exercised, by tactic, with the outcome. This is the boundary between *tested and held* and *never tested*. Techniques not listed here were not run, and their absence from the report is not evidence about them.

TACTIC	TECHNIQUE	NAME	OUTCOME	REF
Reconnaissance	T1595.002	Vulnerability scanning	Broke at detection	B-01
Reconnaissance	T1593	Search open websites and domains	Not applicable	B-05
Reconnaissance	T1598.004	Spearphishing voice	Broke at detection	B-06
Initial access	T1566.002	Spearphishing link	Actioned	B-02
Initial access	T1078.004	Valid accounts: cloud accounts	Held / broke at detection	B-03, B-07
Credential access	T1110.002	Password cracking	Not applicable	B-10, B-14
Credential access	T1110.003	Password spraying	Held at control	B-04
Credential access	T1557.001	LLMNR/NBT-NS poisoning	Broke at telemetry	B-09
Credential access	T1557.001	NTLM relay, SMB and HTTP enrollment	Held at control	B-11, B-36
Credential access	T1557.003	DHCP spoofing	Held (core) / broke (legacy)	B-26, B-27
Credential access	T1558.003	Kerberoasting	Broke at decision	B-13
Credential access	T1558.004	AS-REP roasting	Held at control	B-15
Credential access	T1552.001	Credentials in files	Broke at telemetry	B-24
Credential access	T1556.006	Modify authentication process: MFA	Broke at detection	B-06
Credential access	T1187	Forced authentication	Held at control	B-21, B-28
Credential access	T1003.006	OS credential dumping: DCSync	Broke at decision	B-22
Discovery	T1046	Network service discovery	Broke at telemetry	B-08
Discovery	T1087.002	Account discovery: domain account	Broke at detection	B-12
Discovery	T1135	Network share discovery	Broke at detection	B-23

TACTIC	TECHNIQUE	NAME	OUTCOME	REF
Discovery	T1482	Domain trust discovery	Broke at telemetry	B-35
Privilege escalation	T1649	Certificate template and permission enumeration	Broke at telemetry	B-17
Privilege escalation	T1649	Certificate request with requester-supplied subject	Broke at detection	B-18
Privilege escalation	T1649	Certificate web-enrollment surface	Broke at detection	B-20
Defense evasion	T1550	Use alternate authentication material	Broke at detection	B-19
Persistence	T1671	Cloud application integration	Broke at detection	B-33
Defense evasion	T1550.001	Use of the resulting application access token	Broke at detection	B-33
Execution	No mapping	Business logic abuse of the disbursement API. No ATT&CK technique describes this cleanly; recorded here rather than forced into a category.	Broke at detection	B-31
Lateral movement	T1078.003	Valid accounts: local accounts	Broke at decision	B-16
Persistence	T1053.005	Scheduled task	Broke at decision	B-29
Persistence	T1078.001	Valid accounts: default accounts	Broke at telemetry	B-25
Persistence	T1484.001	Group Policy modification	Broke at telemetry	B-34
Collection	T1213	Data from information repositories	Actioned	B-30

Deliberately not exercised

TECHNIQUE CLASS	WHY NOT, AND WHAT THAT MEANS FOR THIS REPORT
Impact techniques (T1486, T1490, T1489)	Excluded by the rules of engagement. Nothing in this report says anything about the organization's resilience to destructive action.
Exploitation of public-facing applications (T1190)	Attempted against 214 services; nothing exploitable was found. This one is a result, recorded as H-01.
Hardware and physical access (T1200, T1091)	Out of scope by client decision. Untested.
Supply chain compromise (T1195)	Out of scope. Untested and unmodelled.
Endpoint defence impairment (T1562)	Not run. The endpoint estate was well instrumented throughout and we had no operational need to test evasion, so no claim is made about how the endpoint platform behaves under tampering.

D APPENDIX

D Validation conditions

Each entry states the condition that counts as fixed, in terms that can be proven false. "Patch applied" is not a condition; "the same behaviour no longer produces the same outcome" is.

These are validation conditions, not runbooks. They tell your team what has to be true, not the exact commands, tooling, safety guardrails, and cleanup steps needed to establish it. A companion runbook per condition is supplied with the engagement and is deliberately not published in a specimen.

ID	FINDING	WHAT TO RUN	PASSES WHEN
TC-01	F-01	An authorized simulated service desk call requesting a factor reset, using only publicly available information about the target.	The request is refused and escalated; and a legitimate factor change appears in the SIEM within 15 minutes, correlated to any following unfamiliar-device sign-in.
TC-02	F-02	Request a non-existent name from a host on `10.90.0.0/16`; run a listener for four hours; attempt to reach the segment from the virtual desktop estate.	Only the authoritative DNS server responds; the listener collects nothing; the segment is unreachable from the virtual desktop estate.
TC-03	F-03	Enumerate all accounts carrying service principal names; request tickets for each; run the 72-hour offline-recovery companion runbook after recording hardware, software, candidate-set hashes, mask spaces, phase counts, and cost.	Every such account uses a managed service account or has a password set after remediation of 25 characters or more; no password is recovered in the declared run; `svc-bkp-legacy` holds no Tier-0 membership.
TC-04	F-04	As a principal in `NWM-App-Service-Accounts`, attempt to obtain a certificate naming another identity from every template in the forest.	No template issues one; every certificate issued produces a SIEM event within 15 minutes; certificate-based authentication records the issuing template and serial.
TC-05	F-05	Using the integration identity, attempt to move a synthetic disbursement to the approved state through the API.	The transition is refused without a distinct, entitled, human approver; every transition log line carries initiator, approver, and channel; an approval by a service identity alerts within 15 minutes.
TC-06	F-06	Search all indexed shares for the workbook and for credential patterns; check credential set dates; perform a control bulk read of `Finance$`.	No copy exists; all 47 credentials show a set date after remediation; the scan returns no hits above the agreed threshold; the bulk read alerts within 15 minutes.
TC-07	F-07	Attempt the vendor default credentials on the appliance from a corporate desktop; run the estate-wide default credential sweep.	Defaults are rejected; the interface is unreachable from a corporate desktop; an administrative sign-in reaches the SIEM within 15 minutes; the sweep reports zero devices accepting documented defaults.
TC-08	F-08	Attempt authentication with the previously recovered local administrator credential across all 43 hosts, then across a sample of 20 others.	Authentication fails everywhere; every workstation holds a unique managed password with no exceptions; the resulting alert is actioned or closed with a verified reason within the shift.
TC-09	F-09	Attempt to set a password matching the observed pattern; run a repeat spray against a sample of 200 accounts.	The password is rejected at the point of setting; scheduled rotation is disabled for user accounts; the spray alerts within five minutes.

ID	FINDING	WHAT TO RUN	PASSES WHEN
TC-10	F-10	Send DHCPv6 ADVERTISE/REPLY offering an operator DNS server, and separately send unsolicited router advertisements, on `10.90.0.0/16`; observe three hosts on different switches.	No host adopts the offered configuration.
TC-11	F-11	Re-run the day 17 and day 20 behaviours during a normal shift; separately sample the preceding month's Tier-0 adjacent alert closures.	Both produce a containing action, or a closure with a structured reason naming a verified source, within the shift; every sampled closure carries a verification note.
TC-13	F-13	Create an inbox rule that marks matching mail read and moves it to another folder in the same mailbox.	The rule creation alerts within 15 minutes with the actor, mailbox, and rule actions.
TC-14	F-14	As a standard user, attempt to consent an application to mail and file access; then use a controlled public-client authorization-code/PKCE grant to test token refresh before and after the account-compromise runbook's explicit session revocation.	Consent is refused and routed to admin approval; the attempt alerts within 15 minutes; token refresh succeeds before and fails after explicit revocation; the record names client type, flow, token class, reset or revocation mechanism, and both results.
TC-12	F-12	Run authenticated discovery and reconcile to the inventory; introduce a test host and observe its propagation.	Under 2% unexplained difference; every application holding regulated data has a named owner; the test host appears in the inventory and the log-onboarding scope within one cycle.

These conditions are the client's to keep, and the runbooks that accompany them are written to be run by an internal team on a cadence rather than only by us at a retest. A finding that can only be re-verified by the firm that found it is a dependency, not a deliverable.

E APPENDIX

E Evidence index

Every executed behaviour in section 08 has at least one evidence item here. Items are held encrypted and are available to the named distribution list on request. Screenshots, packet captures, and log extracts are not reproduced in full in this document because several contain credential material and internal identifiers. Each item carries an SHA-256 hash recorded at capture time so its integrity can be verified independently.

ID	CONTENT	CAPTURED (UTC)	SHA-256 (FIRST 12)	REFERENCED BY
EV-01-02	External service enumeration output, both ranges	2026-03-02 - 03-09	d01e479f9570	H-01
EV-03	Archived copy of the public support article describing factor-reset verification	2026-03-06 11:20	178792d25c0b	F-01
EV-04-05	Phishing campaign configuration and result set	2026-03-10 - 03-12	6af4ba3d8b5d	H-08, F-01
EV-06	Conditional access refusal events for 11 phished credentials	2026-03-11 09:44	f9b55e254e32	H-04
EV-07	Service desk call: operator notes and timeline. *Recording not retained; the agent is not identified in any artefact.*	2026-03-15 14:12	0e6eb2b8dd08	F-01
EV-08	Password spray result set and detection latency measurement	2026-03-15 - 03-16	5260835c5942	F-09, H-04

ID	CONTENT	CAPTURED (UTC)	SHA-256 (FIRST 12)	REFERENCED BY
EV-09–10	Name-resolution capture summary, legacy segment. *Credential material redacted.*	2026-03-17 08:10	4250f067d6f9	F-02
EV-11	Service account recovery summary. *Values withheld; compute environment and candidate-set details were not preserved, so elapsed times are not a reproducible benchmark.*	2026-03-18 – 03-19	f08431424ab9	F-03
EV-12	SIEM alert record and closure reason, day 17	2026-03-18 09:47	b6e1557acbb1	F-11
EV-13–15	Certificate template properties, CA database query by request ID and serial, request, and issued certificate. Retained fields include requester, template, disposition, issued subject/SAN, serial, and issue/expiry times. *Serial withheld from this document.*	2026-03-19 – 03-20	4bfb9b03277a	F-04
EV-16	Lateral authentication record across 43 hosts, with the associated alert and closure	2026-03-25 16:40	d2ba3048c17a	F-08, F-11
EV-17	Monitoring appliance interface capture. *Topology redacted.*	2026-03-26 10:05	dca322691746	F-07
EV-18	Credential register metadata and access control listing. *File contents not retained.*	2026-03-23 14:30	1849242545ba	F-06
EV-19–20	Disbursement API request and state transition, day 31 access-review ticket, and Finance reversal confirmation. *Exact reversal time not retained.*	2026-03-28 15:41	492957253fee	F-05
EV-21	Replication request alert and closure, day 20	2026-03-21 09:16	f98cece2690e	F-11
EV-22	Reconciled OBJ-2 timeline, operator log against defensive observer record	2026-04-03 10:02	c12df933b670	H-05
EV-23–24	Router advertisement results, core and legacy segments	2026-04-07	1c4b1f25aad1	F-10, H-09
EV-25	Inventory reconciliation working set	2026-04-09	5dd106e1d785	F-12
EV-26	One certificate-based authentication and Event 39 from the servicing domain controller. *DC build, last cumulative update, Kerberos audit-policy setting, and Event 4768 result were not captured; no event is claimed on the other three DCs.*	2026-03-20 11:26	cbed78c3e4cc	F-04, B-19
EV-27	Group Policy object modification and forest trust enumeration output	2026-04-09	6631d3206def	F-04, B-34, B-35
EV-28	Scheduled task persistence: task definition, host, and the alert and closure record	2026-04-06 14:12	6f95c36c3349	F-11, B-29

ID	CONTENT	CAPTURED (UTC)	SHA-256 (FIRST 12)	REFERENCED BY
EV-29	Same-mailbox move/read-mark rule plus public-client application consent and token use; includes tenant audit records, authorization-code/PKCE flow, delegated permissions, and token class. *No password-reset survival test was run during the operation.*	2026-04-08	cbf0b1e007cd	F-13, F-14
EV-30	Security-information deletion, Temporary Access Pass registration, new-method registration, and the sign-in that followed	2026-03-15 14:12	d8c3967faa05	F-01, B-06, B-07
EV-31	NTLM relay attempts: SMB signing rejections and one HTTP enrollment rejection. *IIS EPA settings, SPNs, provider order, template, and HTTP/IIS status were not preserved, so the HTTP cause is unverified.*	2026-03-17 20:00	06808a9d915d	H-06, H-10, B-11, B-36
EV-32	Critical-disclosure ledger: possession, verbal notice, written notice, acknowledgement, and resume records. *Missing D19 and D21 fields are marked rather than inferred.*	2026-03-19 – 03-22	b320af0d181a	section 03, OBJ-1, OBJ-3
RV-01–14	Retest evidence, one item per validation condition	2026-06-08 – 06-12	e9af39909da8	section 12

F APPENDIX

F Tooling and glossary

F Tooling

Listed so the client can distinguish what was proven from what was assumed, and so their own team can reproduce the discovery work. Four of these are recommended for defensive use and are marked; running your own attack-path analysis is cheaper than paying someone to run it for you.

PURPOSE	TOOLING CLASS	NOTE
External enumeration	Port and service scanners; certificate transparency and passive DNS sources	Public-source collection only until the engagement window opened.
Web application assessment	Intercepting proxy, manual testing	Applied to the external estate and to CLARA and POLARIS interfaces.
Directory analysis	Attack-path analysis over Active Directory DEFENSIVE USE RECOMMENDED	The same analysis run internally would have surfaced the F-03 and F-04 chain before we did.
Certificate services analysis	Template and enrollment permission enumeration DEFENSIVE USE RECOMMENDED	A quarterly run would have caught F-04 in 2019.
Directory hygiene assessment	Configuration scoring against known misconfiguration classes DEFENSIVE USE RECOMMENDED	Several of these are free.
Network position tooling	Name-resolution responders, relay tooling, IPv6 advertisement tooling	Used only within the declared operator position and only on in-scope segments.
Credential recovery	Offline recovery against GPU compute	The April run omitted the hardware and candidate-set detail required for exact reproduction. TC-03 therefore requires a fully declared 72-hour companion runbook and does not claim equivalence to that run.
Command and control	Commercial and open-source frameworks	Configuration and infrastructure disclosed at close-out; command infrastructure torn down at the end of the operation on 11 April.
Evidence handling	Timestamped operator logging with per-item hashing	Every evidence item carries a capture-time SHA-256 recorded in Appendix E.

Specific tool names and versions are recorded in the engagement log and available to the distribution list on request. They are omitted here because the list dates faster than the report and because the technique, not the tool, is what a defender needs to test against.

G Glossary

TERM	MEANING AS USED IN THIS REPORT
Adversary validation	Offensive testing whose deliverable is evidence about defensive behaviour, not a list of weaknesses. The difference is that a penetration test ends when a path is proven and this ends when the defensive response to that path is recorded.
Adversary simulation	Realistic attacker behaviour selected from the client's critical paths for what it will prove about the defence, without claiming to reproduce a particular group. What this engagement was.

TERM	MEANING AS USED IN THIS REPORT
Adversary emulation	Reproduction of a named threat actor's known tradecraft, in their known sequence, derived from intelligence on that actor. A different engagement, with an intelligence phase in front of it. Not what this was.
Chain verdict	The furthest point a behaviour reached along the evidence chain before it stopped. See Appendix B.
Recorded	The first moment a source system wrote a record of the activity, whether or not anyone was watching. A log line counts. Reported separately from detection throughout.
Time to detect	Interval from an adversary action to the first analytic or human flag: a detection firing, a ticket raised, a person noticing. Left empty where nothing ever put the activity in front of a decision, which is a result rather than a gap in the measurement.
Time to contain	Interval from an adversary action to the first containing action. Where an objective was also reached, the report states which came first, because containment after the objective has prevented nothing.
Defensive observer	A client-side trusted agent who keeps an independent record of what the defense saw, used to reconcile against the operator log. Their record is authoritative on the defensive side.
Deconfliction	The agreed process by which the client can ask, at any hour, whether activity they are seeing is ours. Invoked twice during this engagement.
Evidence chain	Adversary action, control, telemetry, detection, decision, evidence. The notation used throughout this report and on every Offensive Guardian deliverable.
Persistent issue catalogue	Fourteen issues that account for the overwhelming majority of full compromises observed across a decade of assessments. Every engagement states which were tested and what was found.
Path impact	What the adversary achieved or could achieve, rated against the business. Independent of the chain verdict.
Test case	A reproducible behaviour with a falsifiable pass condition, written so the client can re-run it without us.
Tier-0	The set of systems and identities that control the identity plane. Compromise of any Tier-0 asset is compromise of the forest.
Trusted agent	A named person on the client side who knows the engagement is running. Two were designated here.

ABOUT OFFENSIVE GUARDIAN

Founder-led adversary validation. We recreate the attack paths that matter and show you, control by control, what holds, what fails, and what to fix next. Every engagement is delivered by the practitioner who scoped it.

Led by Jean-François Maes: lead author of SANS SEC565 Red Team Operations and Adversary Emulation, co-author of SEC699 Advanced Purple Teaming, SANS Certified Instructor, former Cobalt Strike researcher at Fortra.

STANDING NOTICE

No engagement, and no report, guarantees that a system is secure or that an intrusion will not occur. This document records what was tested, what was observed, and what was verified at retest. It is an input to risk decisions and not a substitute for them.

Offensive Guardian disclaims liability for damages arising from reliance on this document beyond the terms of the engagement contract.

Evidence, not assumptions.

OFFENSIVE-GUARDIAN.COM