CONFIDENTIAL · PREPARED FOR ACME CAPITAL

INCIDENT RETROSPECTIVE

Initial Access Retrospective — April 2026

11 April 2026. Targeted phishing ’captured Entra session ’OAuth-app persistence ’federated AWS admin ’Secrets Manager read ’12 GB S3 staging and weekend exfiltration. Detected Monday morning by a customer support ticket reporting unusual outbound traffic.

DEMO — NOT YOUR TENANT DATA

This document was generated from CyberTwin's Acme Capital reference fixture so prospects can evaluate the artefact format. It is not derived from your environment. Do not circulate as if it were your own tenant's assessment.

ORGANIZATION

Acme Capital

Reported loss $1.84M · 8 kill-chain phases observed

COUNTERFACTUAL AVOIDED LOSS (MODELED MIDPOINT)

$1.47M

Headline finding

Closing the MFA exclusion on CloudOps would have terminated this incident at minute 14 — saving an estimated $1.47M in breach response, regulatory fines, and customer notification cost.

Kill-chain progression

Initial Access then Persistence then Privilege Escalation then Credential Access then Discovery then Collection then Exfiltration then Impact

Earliest break-point: Initial Access

One open finding at incident time would have stopped the chain at this stage. Remediating any one of them before the attack window would have broken the chain here — averting a modeled share of the incident's loss (breaking earlier avoids more). The avoided figure is a modeled estimate, not a measured amount.

· [CRITICAL] CloudOps group excluded from MFA Conditional Access policy

Detected: 4/13/2026, 8:42:00 AM · Contained: 4/13/2026, 3:10:00 PM

Kill-chain phase analysis

Each MITRE ATT&CK kill-chain phase, with the findings open at incident time that would have prevented or detected that phase.

Initial Access

OBSERVED IN INCIDENT

A targeted phishing email landed on a CloudOps engineer. The Conditional Access policy excluded the CloudOps group — the captured session token was replayed without an MFA challenge. Closing the exclusion alone would have terminated the attack at minute 14.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [CRITICAL] CloudOps group excluded from MFA Conditional Access policy

Execution

No execution payload required — the attacker used the captured Entra session as a legitimate user.

Persistence

OBSERVED IN INCIDENT

The attacker created an OAuth application and self-consented to a refresh-token scope, giving them durable access independent of the captured session.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [HIGH] User consent for low-permission OAuth apps allowed by default

Privilege Escalation

OBSERVED IN INCIDENT

Federated SSO into AWS landed the attacker on the CloudOps engineer's admin role. Console MFA would have produced a second authentication boundary.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [CRITICAL] cloudops-admin AWS IAM user lacks console MFA

Defense Evasion

No active evasion observed; the attacker operated as a legitimate signed-in user.

Credential Access

OBSERVED IN INCIDENT

The admin role pulled the production database master credential out of Secrets Manager — a per-secret IAM scope would have required a separate, alertable elevation.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [HIGH] CloudOps role has unrestricted Secrets Manager read

Discovery

OBSERVED IN INCIDENT

CloudTrail captured 412 lookups across S3, RDS, and EC2 over a 6-hour window. Without lookup-rate alerting, this pattern was invisible until post-incident analysis.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [MEDIUM] CloudTrail lookups not alerted

Lateral Movement

The attack stayed inside AWS — no Defender-monitored Windows endpoints were touched.

Collection

OBSERVED IN INCIDENT

A 12 GB customer-data export was staged in a temporarily-public S3 prefix. Bucket-level Block Public Access would have rejected the policy change at the API.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [HIGH] Sensitive bucket lacks Block Public Access at the bucket level

Exfiltration

OBSERVED IN INCIDENT

A weekend export of 12 GB triggered no alert. GuardDuty in the production account would have paged the on-call within minutes of the first 100 MB.

FINDINGS THAT WOULD HAVE BROKEN THIS PHASE

· [HIGH] GuardDuty anomalous-egress detector not enabled in the production account

Impact

OBSERVED IN INCIDENT

Customer notification, regulator filings, and 28 days of incident-response time are the recurring impact phases.

Methodology + sources

This retrospective walks the MITRE ATT&CK kill-chain phases observed in the incident and lists, for each phase, the findings that were already open in your environment at the time and would have broken the chain if remediated. Findings remediated before the incident date are excluded.

Phase mapping

Phases follow the MITRE ATT&CK Enterprise tactic taxonomy. Each finding's MITRE-tactic tag determines which phase(s) it can break.

Counterfactual loss calculation

The loss avoided is a MODELED range, not a fixed multiplier. Each kill-chain phase carries a 'share of loss already committed by the time it runs' — small at initial access, near-total by impact — so breaking the chain EARLIER avoids strictly more: avoided = 1 − committed at the break phase. A ± band around that midpoint reflects modeled uncertainty, narrowing as more and higher-severity findings corroborate that the break would have held. Every input — the break phase, its committed-loss fraction, the corroborating findings, and the resulting range — is stated in this incident's basis line so an auditor can re-derive it. This is a modeled counterfactual, not a measured figure.

What this is not

This is not a forensics report. We don't reconstruct the attacker's actions; we evaluate which of your already-known findings, if remediated, would have changed the outcome.

DOCUMENT HISTORY

VERSION	GENERATED	ENGINE	ENVIRONMENT	CHANGES
–	5/7/2026, 12:00:00 PM	`cybertwin-engine 0.1.0`	`acme-capital`	Generated

FULL METHODOLOGY

cybertwin.io/methodology — the full catalog of entries, calibration sources, and cross-vendor stitching rules behind the numbers in this report.

SIGNATURES

Prepared by CyberTwin
CyberTwin Engine
Date: 5/7/2026
Counterfactual computed from open-finding remediation hypothesis.

Incident response lead (customer)
Date: _____
Sign acknowledging the retrospective findings and committing to break-point remediation timelines.

Incident Retrospective prepared by CyberTwin · Generated 5/7/2026, 12:00:00 PM

APPENDIX · GLOSSARY

Glossary of terms

Definitions for technical terms used throughout this report. Citations link to the source authority for each definition where applicable.

ALE	Annualized Loss Expectancy. The expected dollar loss per year from a given threat, computed as breach probability × impact. Reported as p10 / p50 / p90 (best / median / worst case) to make the variance honest. *NIST SP 800-30 Rev 1 risk-assessment methodology.*
EDR	Endpoint Detection and Response. Defender for Endpoint, CrowdStrike Falcon, SentinelOne, etc. Distinguished from "AV" by behavioral detection + investigation tooling beyond signature-based malware blocking.
MFA	Multi-Factor Authentication. CyberTwin distinguishes basic MFA (SMS, push) from "phishing-resistant MFA" (FIDO2 / WebAuthn / smartcards) — the latter satisfies stricter framework controls.
MITRE ATT&CK	Adversary tactics + techniques knowledge base maintained by MITRE. CyberTwin annotates every attack-path edge with a tactic code (TA0001 Initial Access, TA0008 Lateral Movement, etc.) and a technique code (T1078.004 Valid Accounts: Cloud Accounts, T1190 Exploit Public-Facing Application, etc.), sourced from the v15.1 Enterprise matrix. Every technique cited in this report traces to a specific edge — the codes are not a fixed count. *attack.mitre.org/matrices/enterprise*
T-code	A MITRE ATT&CK technique code, e.g., T1078.004 Valid Accounts: Cloud Accounts, T1190 Exploit Public-Facing Application. CyberTwin annotates every attack-path edge with a technique sourced from the v15+ Enterprise matrix.
TA-code	A MITRE ATT&CK tactic code, e.g., TA0001 Initial Access, TA0008 Lateral Movement, TA0010 Exfiltration. CyberTwin annotates every attack-path edge with a tactic.

APPENDIX · BIBLIOGRAPHY

Sources cited

Authoritative sources backing the claims in this report. Primary-source documents, published benchmarks, and vendor-cited references. Inline-cited sources throughout the body trace back to entries here.

CISA Known Exploited Vulnerabilities Catalog
KEV — vulnerabilities with confirmed in-the-wild exploitation
https://www.cisa.gov/known-exploited-vulnerabilities-catalog
Accessed 2026-04-01

Lockheed Martin Cyber Kill Chain
Reconnaissance > Weaponization > Delivery > Exploitation > Installation > Command & Control > Actions on Objectives
https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html
Accessed 2026-04-01

Mandiant M-Trends 2024
Dwell time, initial access vectors, and threat-actor TTPs
https://www.mandiant.com/m-trends
Accessed 2026-04-01

MITRE ATT&CK Enterprise Matrix
v15.1 (2024) — tactic and technique IDs cited per edge
https://attack.mitre.org/matrices/enterprise/
Accessed 2026-04-01

Verizon Data Breach Investigations Report 2024
Industry incident pattern and breach-vector statistics
https://www.verizon.com/business/resources/reports/dbir/
Accessed 2026-04-01