paritybit security>

CASE STUDY
Incident Response & Managed Detection
Anonymous Client Engagement

From First Alert to Full Attribution

How ParityBit Security contained a targeted breach against a mid-sized Canadian firm — and rebuilt their human defences in the 90 days that followed.

Client name, sector specifics, and identifying details have been withheld at the client's request under a mutual non-disclosure agreement. All figures are drawn from the engagement record.

At a Glance

Client	Mid-sized professional services firm, Canada
Sector	Legal / advisory (anonymised)
Engagement Period	Managed detection since October 2025; incident response activated February 2026
Platforms In Play	Vector SIEM • Threat Atlas • ZeroPhish
Time to Containment	Under 72 hours from first alert
Material Data Loss	None
Regulatory Breach Notifications	None triggered
Phishing Resilience	Click rate reduced from 22% to under 5% in one quarter

"ParityBit didn't just stop the attack. They told us who was behind it, how they did it, and how to make sure it doesn't happen again.

That is the difference between a vendor and a partner."

— Chief Information Security Officer, Client firm

The Situation

Our client is a mid-sized professional services firm headquartered in Canada. Like many firms in their sector, they sit on a rich vein of confidential client information; and like many firms their size, they had invested in good foundational controls but lacked the 24/7 detection capability and threat intelligence depth of a much larger enterprise.

That gap was the reason they had brought ParityBit on in the first place. Since **October 2025**, their environment had been continuously monitored by **Vector**, our AI-powered unified SOC platform, as part of a managed detection engagement. Telemetry from identity, endpoint, cloud, email, and network layers flowed into a single correlated view, with our analysts on watch around the clock.

On the morning of **Friday, 6 February 2026**, that investment earned its keep. Vector's correlation engine raised a cluster of high-severity alerts against the Finance team's cloud tenant: improbable-travel sign-ins, a short chain of mailbox forwarding rules that no one on staff had configured, and a quiet sequence of outbound transfers timed to avoid the client's legacy thresholds. By the end of that afternoon the incident had been formally activated. By Monday, **9 February 2026**, we knew what had happened, who was behind it, and what to do next.

How We Found It — Vector SIEM

The most important thing to understand about this engagement is that nothing here was a scramble. We did not arrive on-site, install a SIEM, and hope to see something. **Vector had already been watching** for four months. The incident was not discovered because someone got lucky. It was discovered because the architecture did exactly what it was designed to do.

Vector is a multi-layer detection platform. Each layer does a different job, and no single layer alone would have caught this attack. Together, they did.

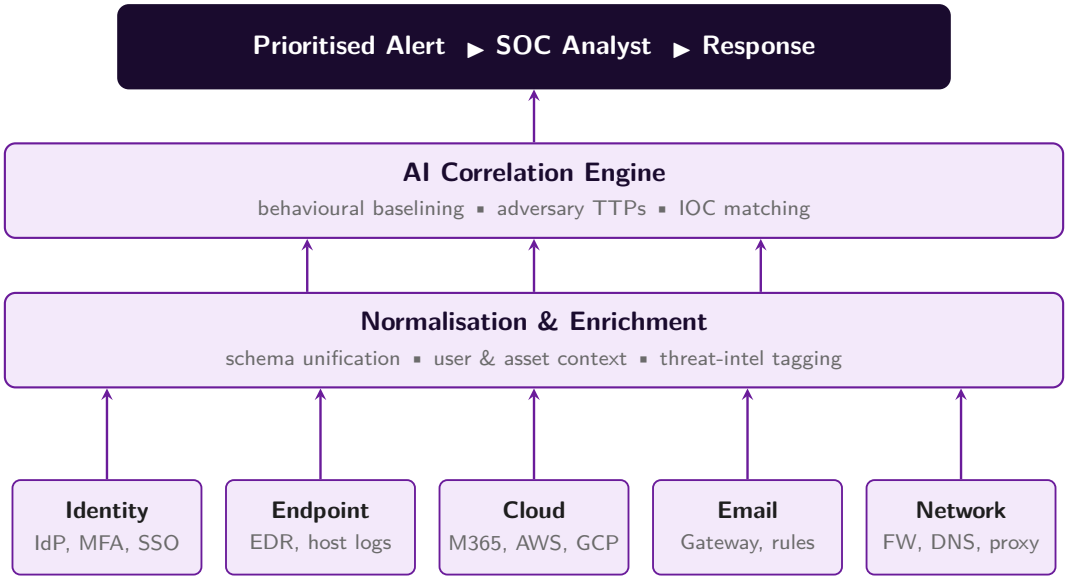

Vector SIEM's multi-layer detection pipeline — the architecture that raised the alert on 6 February 2026.

At the bottom of the stack, Vector continuously ingests telemetry from every layer of the client's environment — identity, endpoint, cloud, email, and network. In the middle, a normalisation and enrichment layer turns that raw data into a common language and adds context: which user, which asset, which known indicator. Above that, an AI correlation engine learns what normal looks like for this specific client and compares every new event against behavioural baselines, adversary tactics, and real-time indicators of compromise. At the top, only the signal that matters reaches a human analyst.

On 6 February, it was the combination that did the work. A single impossible-travel sign-in would have been ignored by most tools. A single mailbox rule would have been missed. A single outbound transfer would have slipped under the threshold. Vector saw all three happen to the same user, in the same hour, on the same tenant — and raised a single, prioritised alert with the full chain already assembled.

That is what "multi-layer detection" actually means in practice. Not a dashboard full of green ticks. A story, told in one alert, about an attacker who had been patient.

What We Learned — Threat Atlas

Detecting the attack was the beginning. Understanding it was the next step — and understanding, for a boardroom, means being specific. "We saw a phishing campaign" is not an answer anyone can act on. *This* is:

As soon as Vector had surfaced the chain of indicators, we pushed them into **Threat Atlas**, our AI-driven threat intelligence platform. Threat Atlas does not simply catalogue indicators; it correlates them across an APT and threat-actor attribution engine, real-time dark web surveillance, and an industry-specific attack matrix. In **3 hours and 20 minutes**, it returned a structured intelligence package.

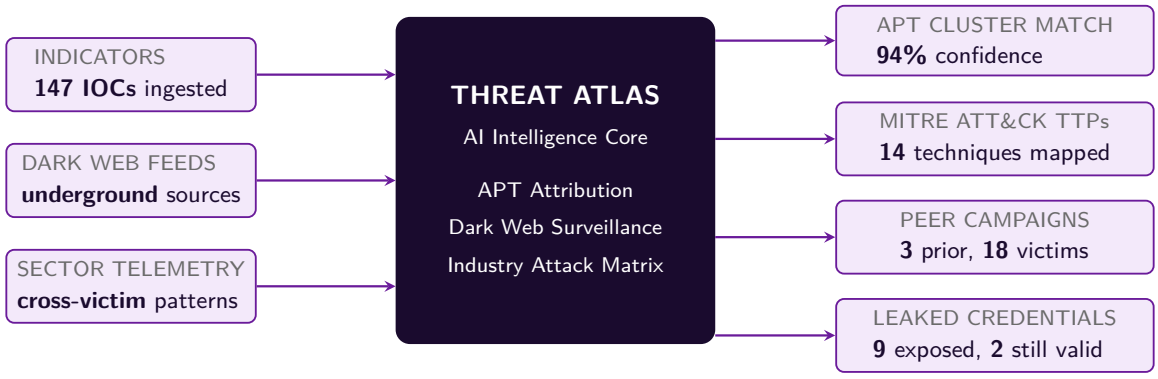

Threat Atlas correlated the incident's indicators against its own intelligence graph — returning an attribution-ready package in under four hours.

Sample Indicators (Representative)

```
adobesign[.]pl                  AiTM phishing kit --- e-signature lure
sapconcur[.]sa[.]com            EvilProxy infrastructure --- expense-report lure
sharepoint[.]za[.]com           Credential-harvest landing page
o365-*-mfa-*[.]{pl,cv,sa.com}   Domain pattern --- MFA-bypass reverse proxy
```

Defanged; cluster & lure infrastructure match publicly-tracked EvilProxy / AiTM phishing operations reported by Okta Threat Intelligence and Microsoft.

The numbers re-framed the conversation for the client's leadership. This was not a freak event:

- **94% of the 147 indicators** matched a single financially motivated cluster Atlas had been tracking since late 2024.
- **14 MITRE ATT&CK techniques** were identified across initial access, credential access, collection, and exfiltration — all consistent with the same operator profile.
- **3 prior campaigns** against peer firms in the same sector matched the pattern, with **18 confirmed victim organisations** across Canada and the US Northeast over the preceding 14 months.
- **9 leaked credential sets** tied to the client's domain were surfaced from underground markets — of which **2 were still valid** at the time of the breach, the probable initial access vector.

It was the last two bullets that mattered most. The client had been partially exposed on the dark web *before* the attack itself began, and the operator behind it was one that ParityBit already knew by name. That reframing gave the client's legal and cyber-insurance teams exactly what they needed: attributable, evidenced intelligence, delivered in a structured report.

For a leadership team who had never before had a named adversary, that clarity was a turning point. It moved the conversation from panic to plan.

How We Hardened Them — ZeroPhish

Every forensic trail led back to the same first move: a single well-crafted phishing email, opened by a single employee. The technical controls had done their job once Vector was in place — the missing layer was human. We deployed **ZeroPhish**, our advanced phishing simulation and awareness platform, across the entire organisation, running every engagement through the same four-stage cycle — *plan, deliver, coach, measure* — so that every simulated attack produces both training and telemetry:

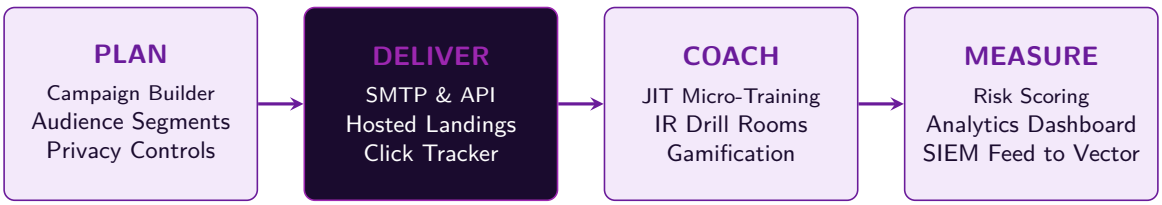

ZeroPhish's four-stage cycle — each simulated attack produces both a training moment and measurable risk telemetry that feeds Vector SIEM.

The first campaign, deliberately modelled on the actual lure that had breached them, produced a **22%**

click rate — a number that quieted the boardroom. Three months later, after targeted simulations paired with integrated micro-training, the click rate had fallen to **under 5%**, and employee reporting of suspicious emails had climbed from **9%** to over **60%**.

The client now runs ZeroPhish campaigns quarterly, paired with incident response tabletop exercises. Muscle memory, not just policy.

The Platform Play

Vector SIEM saw the attack — in minutes, not months.
Threat Atlas named the adversary and surfaced the exposure that enabled them.
ZeroPhish made sure the same first move could not succeed a second time.

Detection. Intelligence. Resilience. One vendor. One timeline.

Outcomes

- Breach contained in under **72 hours** from first alert.
- **Zero** material data loss; no regulatory breach notifications triggered.
- Full adversary attribution delivered to legal, insurance, and (where appropriate) law enforcement.
- Ongoing 24/7 managed detection via Vector SIEM, with mean-time-to-detect reduced from *days* to *minutes*.
- Phishing resilience improved more than **fourfold** in the first quarter; reporting rate improved more than **sixfold**.
- ISO 27001 and PCI audit posture preserved without remediation delay.

Why It Mattered

The question we are most often asked by boards and executive teams is not a technical one. It is simply this: *"If something goes wrong, will we actually be able to handle it?"*

For this client, the answer turned out to be yes — but only because the right platforms, the right intelligence, and the right partner were in place *before* the incident, not after.

That is what we mean when we say **Cybersecurity That Actually Works**.

paritybit security>

About ParityBit Security

A next-generation cybersecurity firm headquartered in Winnipeg, Canada. We build and operate AI-driven platforms for modern defence — **Vector SIEM** (unified SOC), **Threat Atlas** (threat intelligence), and **ZeroPhish** (awareness & simulation) — all aligned to ISO 27001 and NIST, and operated by an in-house team holding OSCP, CISSP, CISM, CEH, CHFI, CompTIA PenTest+, Security+, eCTHP, eCIR, and Certified Red Team Professional credentials.

contact@paritybitsecurity.com • **+1 (204) 963-7230** • **www.paritybitsecurity.com**

Security Without Compromise. Securing your data Bit by Bit.