

Hanco SOC Insights

January–September 2026 | November threat update

Q1 • Q2 • Q3

An anonymised review of identity threats, external exposure and security operations

Prepared by Hanco Global Solutions Ltd

Updated 26 November 2026

From monitoring evidence to practical decisions

This review brings together selected client reporting and operational records to explain what the findings mean for organisations assessing their own security posture. It focuses on actions that can reduce exposure and improve visibility.

Identity	External exposure	Analytical focus
Repeated targeting of a small number of mailboxes	Lookalike domains and service-level vulnerability matches	Evidence quality, business relevance and proportionate action

Client identities and all third-party technology/provider names have been removed. Email addresses, domains, IP addresses and ticket identifiers are also withheld. Figures are scoped to the datasets identified in each section.

Contents

01 Executive summary....3
02 The period in perspective....4
03 Alarm patterns and analytical context....5
04 Identity and phishing....6
05 Brand and external exposure....7
06 Vulnerability and credential exposure....8
07 Connecting prevention and monitoring....9
08 A practical 90-day plan....10
09 Methodology and references....11
10 November security-alert trends....12
11 What the changing mix means....13
12 Industry comparison: threats and exposure....14
13 Industry comparison: SOC priorities....15

How to use this review

Executives can start with the summary and 90-day plan. Security and IT teams can use the evidence sections to shape identity, exposure and monitoring checks. Sections 12–13 compare Hanco observations with current industry research.

Reading the figures

- Reported alarms describe platform detections; they are not counts of successful attacks.
- Phishing-target counts describe targeting recorded in the supplied reports; they are not unique messages or affected people.
- External-exposure findings show where impersonation and exploitation opportunities may exist; prioritise the assets with usable infrastructure and high-severity matches.

01 Executive summary

The evidence points to three priorities: protect repeatedly targeted identities, reduce exploitable exposure and focus SOC attention on meaningful changes in the detection mix. Sections 12–13 compare these findings with current industry research and translate them into operational priorities.

Identity protection needs focus.

In September, two anonymised clients recorded 95 and 106 phishing-target counts respectively. For one client, two mailboxes accounted for 93 of 106 counts (87.7%). Focus protection and user support where targeting is concentrated. These are detections, not confirmed compromised accounts. [E3]

External exposure creates actionable priorities.

September reporting identified 119 lookalike-domain observations: 18 (15.1%) with website infrastructure and three (2.5%) with mail infrastructure. These are the most actionable leads because they could support brand impersonation or deceptive email. Separately, 39 of 60 service/version matches (65%) were High or Critical and concentrated on one asset, making that asset a priority for exposure and patch review. [E4]

Severity is rising within stable alarm volume.

The April–September external-monitoring series recorded 6,154 alarms. Although Q3 volume was almost unchanged from Q2 (+0.2%), High-severity alarms rose 18.2%, from 660 to 780; their share increased from 21.5% to 25.3%. Deep/dark-web and brand-related categories contributed 88.4% of the six-month total. Prioritise severity and exposure, not headline volume. [E1]

A separate late-October/November snapshot recorded 19,234 deduplicated alerts across four environments, of which 82.6% were High severity. Three rules accounted for 78.1% of records. In four equal weeks, total alerts fell 63.7% from the first to the last week, while all other rules increased 52.3%. This points to a changing detection mix, not a uniform reduction in security exposure. [E6]

Decisions for the next 90 days

Assign owners to repeatedly targeted accounts and priority exposed services. Verify identity, endpoint and telemetry controls. Measure detection quality and verified remediation, not just alert counts; industry research supports these priorities (see sections 12–13).

Scope: January–September external-monitoring and operational evidence; separate 26 October–25 November security-alert snapshot. Counts describe detections, not verified incidents. The two periods and monitoring layers are not combined.

02 The period in perspective

Q1–Q3 covers 1 January–30 September 2026. The reconciled April–September series provides a six-month baseline, complemented by early-period identity and operational observations. The November update is analysed separately in sections 10–11.

Period	What the evidence supports	Management implication
Q1 | Jan–Mar	February reports describe phishing and credential-related themes. February/March external threat-monitoring summaries contain conflicting totals and are not used in the trend series.	Treat the early-period findings as directional context; avoid a Q1-to-Q3 growth claim.
Q2 | Apr–Jun	One client's monthly external threat-monitoring category and severity totals reconcile: 971, 1,018 and 1,085 alarms.	Prioritise recurring exposure themes and business relevance within the alarm volume.
Q3 | Jul–Sep	The same dataset reports 1,139, 987 and 954 alarms. September adds detailed phishing, lookalike-domain and vulnerability evidence from other anonymised datasets.	Focus on concentrated identity targeting and validate the exposed services highlighted by the September findings.

Four sources of visibility

Security monitoring provides telemetry and email-security reporting. Endpoint prevention provides a protective layer. External threat intelligence supplies digital-risk and asset observations. Dark web monitoring covers exposed data and related external-monitoring categories. Their measures are complementary and are not added into a single "threats stopped" figure.

What the evidence means for management

Prioritise identity protection, exposed services and monitoring health. The six-month severity shift and September concentration findings give leadership specific control priorities; the November update shows how detection mix and recurring sources affect SOC workload.

Evidence: E1–E5. Quarter boundaries follow calendar quarters. Report dates and export snapshots are distinct from the dates of the underlying observations.

03 Alarm patterns and analytical context

A selected April–September dataset provides the most consistent monthly alarm series. Category totals and severity totals reconcile to each month's reported total. Across the six months, 6,154 alarms are reported; this is a sum of monthly records, not a deduplicated case count. [E1]

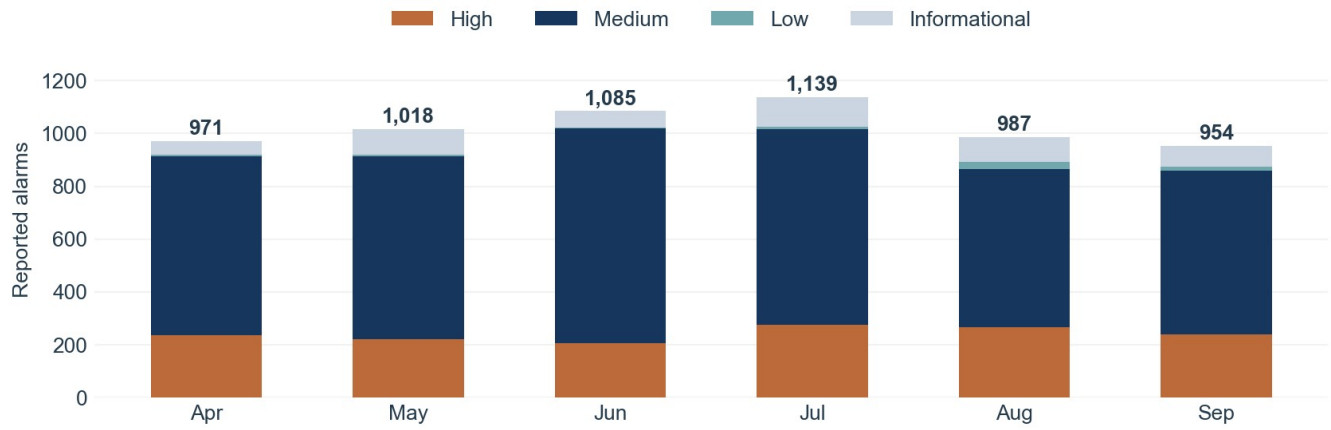


Figure 1. Selected April–September dataset only. High includes Critical where the source report combines these ratings. No inference about confirmed compromise is made from severity.

Measure	Q2	Q3	Apr–Sep
Reported alarms	3,074	3,080	6,154
High (source-normalised)	660	780	1,440
Medium	2,190	1,959	4,149
Low	14	53	67
Informational	210	288	498

Where the alarm volume sits

Deep/dark-web categories account for 4,256 alarms (69.2%); brand and lookalike-domain categories account for 1,181 (19.2%). Vulnerability intelligence contributes 377 (6.1%); the remaining categories total 340 (5.5%). April's separate phishing-domain category is grouped with brand-related findings for comparability.

Q3 High-severity alarms increased 18.2% from 660 to 780 even as total alarms rose just 0.2%. High-severity share rose 3.8 percentage points (21.5% to 25.3%). July was the volume peak (1,139); September was 16.2% lower (954). Falling aggregate volume therefore need not imply a less demanding review workload.

04 Identity and phishing

September's email and identity reporting shows a useful pattern: targeting is concentrated within each organisation. One anonymised dataset records 95 phishing-target counts across 19 listed mailboxes; a second records 106 across nine. These datasets are shown separately because detection counts are not a measure of unique people or unique messages. [E3]

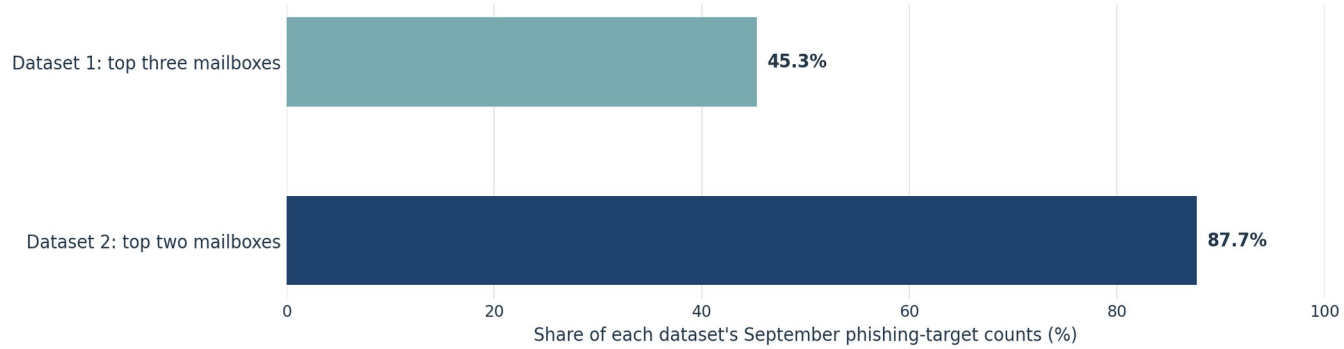


Figure 2. Dataset 1: 43 of 95 counts across the top three mailboxes. Dataset 2: 93 of 106 across the top two. The groups have different sizes and are not a like-for-like organisational ranking.

What the supporting detail adds

- In Dataset 1, the two leading sender IPs account for 80 of 89 sender-IP counts (89.9%). This is a separate measure from its 95 target counts; the two totals are not added together.
- Dataset 1's reported lures include document availability, encrypted-file warnings, missed calls and payment themes. Its leading detection methods include DMARC spoofing and URL reputation.
- Dataset 2's leading reported detection methods include fingerprint matching and brand impersonation. Two mailboxes carry most of the observed targeting burden.

Practical response

Review MFA and access controls for repeatedly targeted accounts, validate email protections, and give users examples based on the lures they actually receive. Correlate email detections with sign-in and account activity before concluding that a mailbox was accessed. UK government cybersecurity identity guidance supports MFA, monitoring and controls against exposed passwords. [G1]

February also showed phishing and authentication anomalies. The repeated appearance of identity-related findings at different points in the year supports sustained identity protection rather than a one-off awareness exercise.

05 Brand and external exposure

The September external-monitoring reports identify 119 lookalike-domain observations across three anonymised clients. This is the sum of client-level distinct-domain counts, not a claim that 119 globally unique malicious domains were found. [E4]

Client	Lookalikes	Website infrastructure	Mail infrastructure
Dataset 1	14	0	0
Dataset 2	75	16	2
Dataset 3	30	2	1
Dataset-level sum	119	18	3

How the names resemble legitimate brands

Top-level-domain substitution accounted for 107 of 119 observations (89.9%). This strong concentration suggests that monitoring alternative domain suffixes is a high-value starting point for brand protection; other naming techniques still warrant coverage.

Which observations deserve attention first

Website infrastructure was observed for 15.1% of lookalikes; mail infrastructure for 2.5%. These are higher-priority leads because they provide a potential route to impersonation or deceptive communications. The groups can overlap. Check authorisation and use before escalation or takedown.

Why this matters to the business

A lookalike with functioning infrastructure creates a more credible route to deception than a name alone. Prioritise web-enabled domains for content review and mail-enabled domains for impersonation and phishing assessment.

Keep asset inventory separate from threat activity

Asset inventories should be reviewed as exposure snapshots. One September source includes early October discoveries; the useful management measure is the number of relevant exposed assets that have been reviewed and resolved.

The observed infrastructure establishes a potential attack surface, not evidence that a malicious campaign has already occurred.

06 Vulnerability and credential exposure

Version intelligence: a priority for validation

One September dataset identified 60 service/version matches on a single asset: eight Critical, 31 High, 19 Medium and two Low. The 65% High/Critical concentration makes the asset a priority for review and remediation. Installed versions, patch backports, enabled features and network reachability determine whether the matched weaknesses are exploitable. [E4]

Evidence type	What was reported	Appropriate next step
Service/version intelligence	60 observations; 39 High/Critical; one affected asset.	Validate installed versions, enabled components, exposure and vendor advisory applicability.
Assessment findings	Dataset 1: five findings across three targets—two Medium and three Low.	Confirm ownership and relevance; test the findings and agree proportionate changes.
Credential exposure	Dataset 2: one record with a September last-seen date and credential-related field types.	Confirm organisational relevance, review affected-account activity and take proportionate credential action.

Assessment detail

A separate assessment recorded two Medium findings involving missing HttpOnly cookie attributes and three Low findings involving TCP timestamp disclosure. Prioritise the cookie configuration where sensitive sessions are involved; the timestamp findings are a lower-priority hardening opportunity.

Credential evidence has a different clock

Credential-related exposure was observed in September. Treat this as an identity-risk lead: confirm account relevance, review recent access, revoke sessions or reset credentials where warranted, and strengthen authentication. The last-seen date is an observation date, not the original breach date.

Risk-based prioritisation

Prioritise the eight Critical and 31 High matches on the affected asset; verify patch level and exposure, then remediate applicable weaknesses and retest. This directs resources to the concentrated risk rather than treating every match as a separate incident. [G2]

07 Connecting prevention and monitoring

The operational worksheet adds context on prevention and monitoring, including device counts, protection observations, source availability and log recency. Its available records extend through 5 August 2026. It supports a view of the checks recorded, not a complete nine-month measure of service performance. [E5]

Layer	Role in this review	Useful question to ask
Endpoint prevention	Prevention, with worksheet fields for devices and protection observations.	Are the intended devices covered, and are coverage exceptions investigated?
Security monitoring	Security telemetry and reporting context, including source availability and recent-log fields.	Are important systems sending usable data, and would a missing feed be noticed?
External intelligence	External threat intelligence, digital-risk and asset observations.	Which findings are relevant to the organisation, and what evidence supports action?
Dark web monitoring	Exposure observations and external-monitoring categories.	Does the observed data relate to the organisation, and has the relevant risk been assessed?

A worksheet snapshot is not an event ledger

Repeated shift readings can describe the same platform state. Use the worksheet as a control-health record: check whether protection coverage and log delivery stay consistent across shifts, assign exceptions to owners and confirm resolution.

Connect the checks to outcomes

Pair protection coverage checks with telemetry tests and a documented route for acting on relevant findings. UK government cybersecurity logging guidance emphasises selecting data for a clear monitoring purpose and using it to investigate suspicious activity. [G3]

Track protection coverage, verified log delivery, repeat detections and remediation closure. The November snapshot further identifies high-volume detection rules and control-tampering indicators that merit targeted SOC review.

08 A practical 90-day plan

The following is a proposed improvement plan from the date of review. It is not a record of actions already performed. Owners and success measures should be agreed against the organisation's priorities and operating constraints.

Window	Recommended focus	Suggested owner	Evidence of progress
Days 1–30	Review repeatedly targeted accounts; validate MFA and email protections. Confirm the most consequential external-service findings.	Identity lead and infrastructure owner	Account/control review recorded; applicability and exposure of priority findings established.
Days 31–60	Review web- and mail-enabled lookalikes for legitimacy. Check prevention coverage and log delivery for critical systems.	Security lead, brand owner and IT operations	Domain decisions documented; coverage exceptions assigned; sample log-delivery tests recorded.
Days 61–90	Retest agreed changes and exercise the path from detection to decision. Establish repeatable monthly measures.	SOC lead and business risk owner	Verification evidence retained; exercise lessons assigned; metrics have clear definitions and denominators.

Questions for a leadership discussion

- Which identities and external services would cause the greatest disruption if misused?
- Can we demonstrate that our prevention and monitoring cover those priorities today?
- Do we know which external findings are relevant, and who owns the next decision?
- Can our reporting show validated improvements without confusing alarm volume with business risk?

A focused starting point with Hanco

A focused discussion with Hanco Cyber can help identify the accounts, internet-facing assets and recurring alert patterns that merit attention. Agree an evidence-based review of exposure, identity controls and monitoring effectiveness, with measurable follow-up actions.

Recommendations respond to the supplied observations. No savings, avoided losses, guaranteed outcomes or prevention-effectiveness percentages are claimed.

09 Methodology and references

Scope, counting and anonymity

Q1–Q3 evidence draws on 15 monitoring reports (98 pages) and a local operational worksheet saved 5 August 2026. The separate November update uses four alert exports. Each analysis retains its own reporting period and unit of measurement.

Monthly external-monitoring alarms are summed within the reconciled April–September series. September reports cover different monitoring layers. Percentages use the stated denominators, and monitoring observations are distinguished from verified incidents.

February/March category totals were not reconcilable and are used as qualitative identity and exposure context. April phishing-domain alarms are grouped with brand-related findings. Observation dates and inventory snapshots remain distinct.

The circulation version omits client and technology-provider names, sectors, personal identities, domains, IP addresses, ticket identifiers, source filenames and credential values. Examples are distinguished only by neutral dataset labels where comparison requires it. No original evidence images or embedded source files are included. Detailed provenance is held separately in internal working files.

Evidence references

Ref.	Source description and use
E1	Selected April–September external-monitoring series; monthly category and severity tables.
E2	Selected February/March external-monitoring reports and February email/identity reports. Early-period context only.
E3	Selected September email and identity reports; target, sender and detection-method tables.
E4	Selected September external-monitoring reports; lookalikes, infrastructure, vulnerability and exposure observations.
E5	Local Hanco SOC worksheet, saved 5 August 2026; operational context only.

Supporting guidance

G1: UK government cybersecurity, 10 Steps to Cyber Security — Identity and access management. G2: UK government cybersecurity, 10 Steps to Cyber Security — Vulnerability management. G3: UK government cybersecurity, 10 Steps to Cyber Security — Logging and monitoring. Guidance consulted 25 November 2026; used for recommendations, not as evidence of client outcomes.

E6: Four anonymised security-alert exports, approximately 26 October–25 November 2026 UTC. Exact duplicate rows within each export are removed. Equal-week comparisons use 28 October–24 November, avoiding partial boundary days. Source severity is retained; no status-code interpretation is assumed.

10 November security-alert trends

26 October–25 November 2026 | Separate post-Q3 monitoring snapshot

20,055 raw records | 19,234 deduplicated alerts | 15,888 High severity | 64 distinct alert names

Four anonymised environments generated 19,234 records after removal of 821 identical rows (4.1% of the export). High-severity records comprised 82.6%. These are detection records, not a count of attacks or incidents.

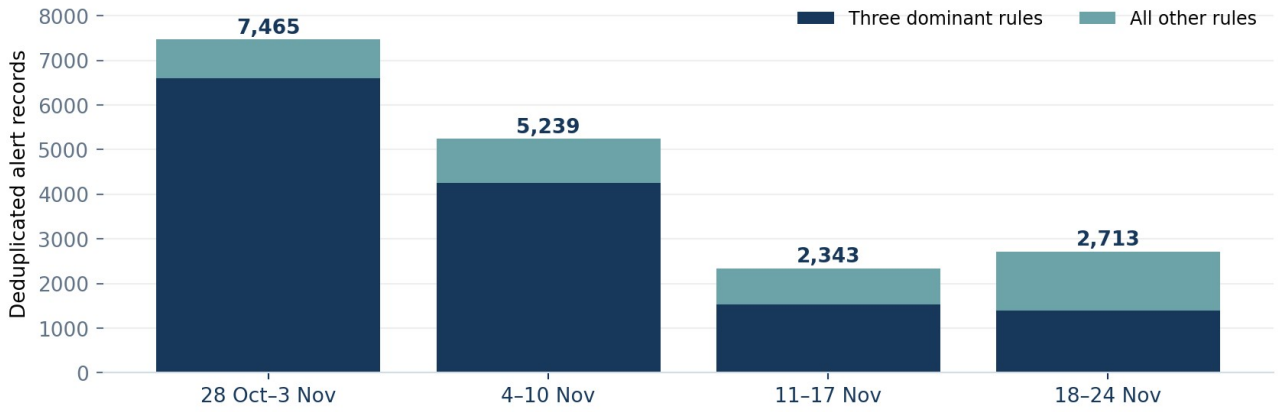


Figure 3. Equal seven-day windows (UTC), excluding partial boundary days. Three dominant rules: suspicious scripting, process masquerading and known-malicious-IP detections.

A falling total masks a changing workload

The weekly count fell from 7,465 to 2,713 (-63.7%). The first two weeks recorded 12,704 alerts versus 5,056 in the final two (-60.2%). But the final week rebounded 15.8% from week three, and records outside the three leading rules rose from 865 to 1,317 (+52.3%) between weeks one and four.

Suspicious scripting fell from 3,879 to 223 per week (-94.3%); meanwhile, login-item persistence rose from 14 to 277 and process masquerading rebounded from 275 to 1,011 in the last week. The alert mix is changing materially, so triage priorities should change with it.

Management implication

Investigate the decline in scripting alerts against changes in monitoring coverage, rule settings and endpoint activity. Continue to review persistence, masquerading and identity activity even when total alerts fall. Sources generating alerts decreased from 195 in week three to 112 in week four; this is a coverage-check trigger, not proof of lost telemetry.

11 What the changing mix means

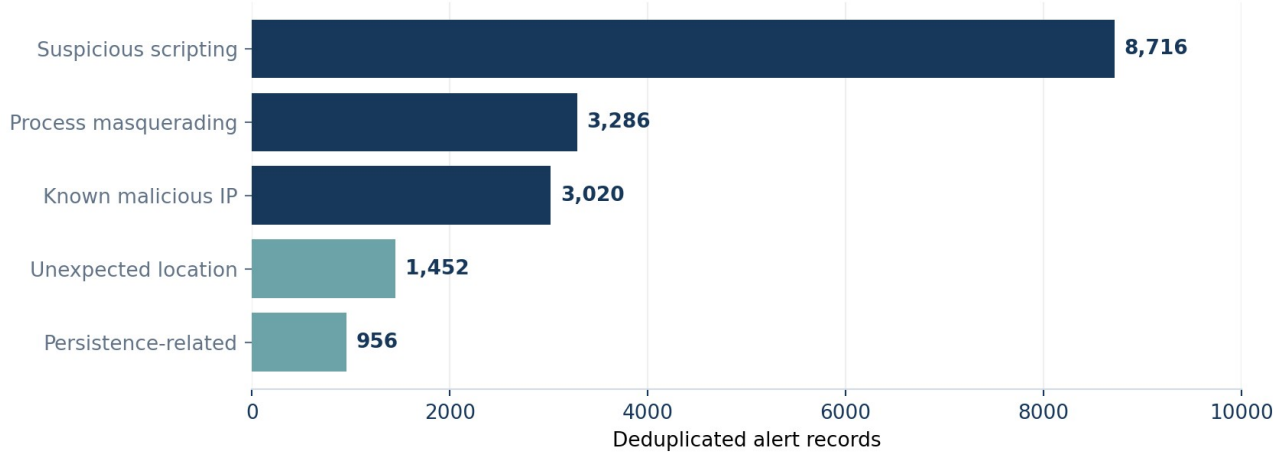


Figure 4. Selected detection categories and grouped themes across the full export window. Counts are records; categories shown are separate, not a complete breakdown.

Concentration is an opportunity for better triage

The three largest rules account for 15,022 records (78.1%). Suspicious scripting appeared across three environments and 173 environment/source pairs. One monitored source generated 3,284 of 3,286 process-masquerading records (99.9%). Review that source and rule context first; repeated patterns may reflect automation or recurring hostile activity, and neither should be assumed without evidence.

Control and identity indicators merit focused review

Unexpected-location activity generated 1,452 records, persistence-related activity 956, external sharing 186, and inbox-rule changes 102. Service principal creation (78), mailbox delegation changes (67) and permission escalation (35) point to privileged identity and collaboration controls as practical review areas. There were 63 firewall or antivirus protection-disabled records and 21 ransomware-indicator records: check control health and correlate with process, user and endpoint evidence.

Turn signals into measurable actions

Prioritise recurring high-severity sources; review approved administrative scripts and changes before tuning rules; verify endpoint protections and privileged access; and track weekly alert mix, repeated-source concentration and verified remediation. Use Hanco Cyber to establish an agreed baseline and a repeatable improvement plan.

12 Industry comparison: threats and exposure

The 2026 industry research provides context for Hanco's observations. External sources describe different populations, time periods and measurement methods; the comparisons below identify common priorities and differences, not equivalent incident rates.

Exploitation and identity: two routes into organisations

Verizon's 2026 DBIR identifies vulnerability exploitation as the leading initial access vector (31% of breaches); Mandiant's 2026 M-Trends likewise reports exploits in 32% of investigated intrusions. Hanco's September dataset had 39 High/Critical service-version matches out of 60, concentrated on one asset. This supports prioritising externally reachable, applicable vulnerabilities and known exploitation evidence before lower-impact findings. [I1, I2; E4]

Microsoft's 2026 Digital Defense Report found that 52.2% of valid-account intrusions involved further credential theft. Hanco observed 87.7% of phishing-target counts concentrated in two mailboxes in one environment, alongside 1,452 unexpected-location alerts and 102 inbox-rule-change alerts in November. These signals support enhanced identity controls for repeatedly targeted or privileged accounts and correlation across email and sign-in activity. [I3; E3, E6]

Differences are as valuable as similarities

Mandiant's 2025 investigations showed voice phishing rising to 11% of initial infection vectors while email phishing fell to 6% from 14% in 2024. Microsoft nevertheless reports identity and impersonation as major threats. Hanco's email-focused findings establish recurring email exposure, but do not measure voice-based attacks. Extend monitoring and exercises to helpdesk impersonation, MFA reset requests and collaboration channels. [I2, I3; E3]

Trusted tools and behaviour-based detection

Red Canary's 2026 Threat Detection Report ranks PowerShell second among observed techniques, while Hanco recorded 8,716 suspicious scripting alerts (45.3% of November records). CrowdStrike reports that 82% of its 2025 detections were malware-free. These observations point to the importance of behavioural detection and context around legitimate tools, not merely malware signatures. Hanco's alert counts do not establish the malware-free proportion or the number of confirmed PowerShell threats. [I4, I5; E6]

Comparison principle: industry breach or threat percentages must not be equated with Hanco rule-trigger percentages. External publications are research citations only; they do not identify Hanco's technology providers.

13 Industry comparison: SOC priorities

A falling headline count can hide emerging priorities

SANS's June 2026 SOC survey found that 70% of respondents still used incident count as their leading SOC metric. Hanco illustrates the weakness of volume-led reporting: November weekly alerts fell 63.7%, while alerts outside the three dominant rules increased 52.3%. In Q3, external-monitoring alarm totals rose only 0.2%, yet High-severity alarms rose 18.2%. Track detection mix, affected sources, triage outcomes and verified remediation alongside volume. [I6; E1, E6]

Detection quality and automation need governance

The 2026 SANS detection engineering study reports that 59% of teams measure false positives but only 14% prioritise their reduction. Hanco's three leading rules accounted for 78.1% of November alerts, and one source produced 99.9% of masquerading alerts. Investigate repeated patterns, group related signals and test rule changes against known detection scenarios. Concentration does not establish a false-positive rate. [I7; E6]

SANS reports that 79% of SOCs use AI/ML, while only 36% have embedded it in defined workflows. CrowdStrike's 2026 report puts average eCrime breakout time at 29 minutes in 2025. Together these findings support governed automation for enrichment and triage, with analyst approval and measured outcomes. Hanco's records do not measure AI usage or incident response speed. [I5, I6]

External research references | 2026 editions

I1 Verizon, 2026 Data Breach Investigations Report

I2 Mandiant / Google Cloud, M-Trends 2026, executive edition

I3 Microsoft, Digital Defense Report 2026

I4 Red Canary, 2026 Threat Detection Report: PowerShell

I5 CrowdStrike, 2026 Global Threat Report

I6 SANS Institute, 2026 SOC Report

I7 SANS / Anvilogic, State of Detection Engineering Report 2026

Sources consulted 26 November 2026. Industry findings provide context, not evidence of incidents within Hanco-monitored environments.