The Annual State of Cyber 2025

Backed by Triskele Labs' reports, revealing the tactics, trends, and takeaways shaping the year ahead.

DFIR GRC MDR Offensive

Behind the Breach: Triskele Labs' 2025 DFIR Findings

Contrary to popular assumptions, most malicious activity did not originate from traditionally expected locations such as Russia, China, or parts of West Africa. Instead, Threat Actors frequently leveraged infrastructure based in more trusted regions to evade suspicion and avoid geo-blocking.

Ransomware and Business Email Compromise dominated, accounting for more than 75% of all cases. Ransomware activity doubled, yet only 4.8% of victims paid a ransom, reflecting improved resilience.

BEC cases surged by 86%, driven by session token theft and persistence within Microsoft 365. DFIR investigations revealed recurring weaknesses in remote access controls, monitoring, and password management, particularly in healthcare and finance.

Read full report

DFIR GRC MDR Offensive

Misguided Advice and Missed Opportunities

Despite positive headlines suggesting progress, Australia's cyber security landscape remains deeply divided. Larger organisations, equipped with resources and expertise, continue to strengthen their defences and influence government policy.

Meanwhile, small- and medium-sized businesses struggle under the weight of compliance requirements and generic advice that fails to reflect operational realities.

As "trickle down" security pressures mount, smaller firms are forced to chase certifications and frameworks without addressing fundamental risks, leaving them vulnerable to opportunistic threat actors.

To close the gap, the focus must shift toward practical, accessible, and contextual cyber advice.

Read full report

DFIR GRC MDR Offensive

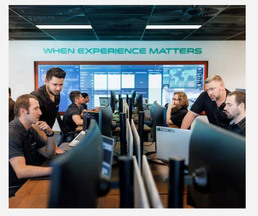

A year of smarter detection, faster response and measurable improvement.

FY25 marked another year of growth and operational excellence for the Triskele Labs Security Operations Centre (SOC). As cyber threats continued to evolve in sophistication and scale, our focus remained on delivering rapid, intelligence-led detection and response across every client environment.

During the year, the SOC processed more than 260,000 alerts — a 40% increase on FY24 — driven by expanding client coverage across finance, health, education, construction and critical infrastructure.

Despite this uplift, alert density per client remained stable, underscoring the efficiency of our tuning, correlation and automation initiatives.

Read full report

DFIR GRC MDR Offensive

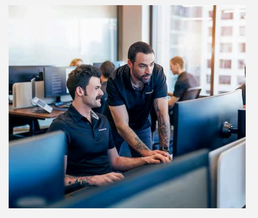

Exposed: What 500+ Pen Tests Reveal

Triskele Labs' State of Cyber — Offensive Report 2025 reveals key shifts in the threat landscape following more than 500 penetration testing engagements conducted across critical sectors during FY2024.

The Offensive Security team identified 3,887 vulnerabilities — averaging eight per engagement—while observing a significant decline in critical findings, indicating the value of ongoing testing and remediation programs.

However, persistent weaknesses remain, with low-severity vulnerabilities rising and recurring flaws in broken object authorisation, authentication, and misconfiguration continuing to dominate results.

Read full report

SPECIAL EDITION OF CYBERS: STATE OF CYBER

Behind the Breach: Triskele Labs' FY25 Findings

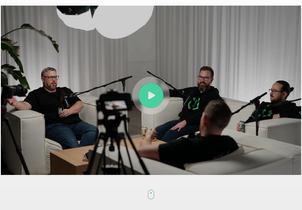

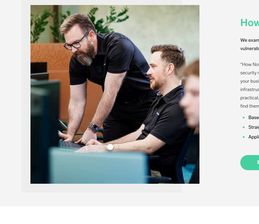

How not to get hacked?

We examine your systems, networks, and processes to uncover vulnerabilities and weak points before attackers do.

"How Not to Get Hacked" is our high value - targeted and quick external security review, designed to uncover what attackers see when they look at your business. We identify exposure across your internet-facing infrastructure, Microsoft 365 configuration, VPN security, and then deliver practical, prioritised advice to close the gaps before threat actors find them.

- Based on real incidents
- Straightforward advice, non-technical language
- Applicable to finance, healthcare, legal & more

Enquire today

Benchmarking Microsoft 365 Security in Australian Organisations

Understanding Your Microsoft 365 Security Posture

Gain a rapid, high-impact assessment of your Microsoft 365 environment. Triskele Labs' Microsoft 365 Baseline Security Review benchmarks your configurations, licensing, and identity management against industry best practice, helping you uncover the weaknesses attackers target most.

The result — a clear, actionable snapshot of your organisation's Microsoft 365 security posture, with practical recommendations to strengthen your defences and align with standards such as the ASD Essential Eight and CIS benchmarks.

Book your M365 Baseline Review

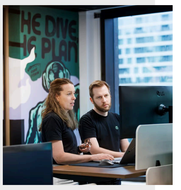

We are the only cybersecurity experts you'll ever need to talk to.

We consider every business we work with a partner — not just a customer. At Triskele Labs, we work with you to understand your risks, goals, challenges and culture to develop Cyber Security solutions tailored to your business. With us by your side, you can be confident that your data and systems are secure — and that you have a trusted, responsive and experienced partner to protect your business.

First name* Last name
Work Email*
Phone number
Any further information
Submit

DFIR

Executive Summary

BEC and Ransomware Dominate the Threat Landscape

Healthcare and Finance Hit the Hardest

Unsecured Remote Access Still a Key Entry Point

Threat Actors are Getting Smarter

Background

Contributors

Incident Statistics

Incident Heatmap

Incident Types

Business Email Compromise

Ransomware

Broader Incident Types

Industry Verticals

66% of malicious activity came from 'trusted' countries

Location alone is not a reliable indicator of malicious activity

Case Study: Healthcare Provider

Initial Compromise via Insecure Remote Access

Privilege Escalation and Unimpeded Data Exfiltration

Sector Specific Challenges

Ransomware Investigations

Overview

Double and Triple Extortion Tactics

Stages of a Typical Ransomware Attack

Ransomware Groups

Group Spotlight: RansomHub

Initial Access

Data Exfiltration

Over the past financial year, we conducted a total of 63 ransomware investigations, with 52 of those cases involving confirmed data exfiltration. The confirmation of exfiltration activity was determined through a combination of network-based forensics, host-based forensics, and ISP metadata, which continue to be critical sources in identifying threat actor behaviour post-compromise. This highlights the ongoing trend where data theft is now a standard component of ransomware operations, not just encryption alone.

Limited Forensic Evidence in Some Cases

Despite strong visibility in many cases, 17 of the 63 engagements did not contain sufficient forensic artefacts to definitively identify the tooling used for exfiltration. This typically occurs when attackers clear logs, perform full disk encryption, or when endpoint or network visibility logging has short retention periods. This underscores the importance of comprehensive logging, monitoring, and endpoint detection to aid incident response investigations.

Preferred Tools: Rclone, Mega, and FileZilla

Among the cases where tooling could be confirmed, Rclone emerged as the most frequently used exfiltration tool, appearing in 13 engagements. Its versatility, command-line interface, and support for a wide range of cloud storage services make it a preferred choice for attackers. Following Rclone, Mega and FileZilla were tied in second place, each used in 7 engagements. These findings reflect a continued reliance on common, off-the-shelf tools by threat actors, likely to evade detection and blend into legitimate activity.

Volume of Data Exfiltrated

The average volume of data exfiltrated per confirmed engagement was 581 GB, representing a decrease from 406 GB observed in the previous financial year. While this drop is notable, it still demonstrates the significant impact of these incidents, particularly where sensitive or regulated data is involved.

Proactive Disruption of Threat Infrastructure

In several engagements, Triskele Labs successfully obtained credentials or configuration files linked to Threat Actor infrastructure, allowing for further action beyond immediate containment. These artefacts were promptly provided to Legal Incident Responders, enabling formal takedown requests to be issued, or were shared directly with Australian and international law enforcement agencies. In multiple instances, this led to the disruption of exfiltration infrastructure, including the successful removal of data repositories where stolen information was being held.

However, even where takedowns have been effective, data is still published in the vast majority of cases where a ransom payment is not negotiated. This reinforces the reality that Threat Actors maintain multiple copies of exfiltrated data, often across different platforms or locations. This was clearly illustrated in recent law enforcement operations targeting LockBit, where data persisted despite claims by the ransomware group that it had been deleted.

Assume Compromise — Always

These findings underscore the importance of assuming data is compromised once exfiltrated, regardless of whether it appears to have been removed, and reinforce the value of timely legal engagement during the incident response process.

Exfil Tooling	Count
Unknown	17
Rclone	13
FileZilla	7
MEGASync	7
WinSCP	4
S3 Bucket	1
Bitvise	1
Cobalt Strike	1
Google Drive	1

Exfiltration tooling observed with incident counts.

Ransom Payments

The Decline of Payouts in 2025

The 2024–25 financial year saw a significant decline in ransom payments, continuing a positive trend in organisational resilience. Out of 63 ransomware investigations conducted, a ransom was paid in only three cases, representing just 4.8% of engagements. This is a marked drop from the 13% payment rate observed in the 2023–24 financial year, indicating both a shift in incident response decisions and growing confidence in recovery strategies.

Several contributing factors have been identified. Firstly, many victim organisations had robust backup and recovery capabilities in place, which enabled them to restore operations without needing to negotiate with Threat Actors. Secondly, a number of organisations were targeted by sanctioned ransomware groups, which made payment legally impossible, regardless of their ability to recover. In other cases, organisations assessed the cost of data recovery, including rebuilding from non-encrypted systems, offsite data, or prior exports and found it to be more economical than paying for a decryption tool. Finally, a growing number of organisations are taking a principled stance, refusing to pay on moral grounds in an effort to disrupt the ransomware business model and avoid reinforcing its profitability.

In addition to the reduced frequency of payments, the average ransom paid also declined. This year, the average payment was approximately $77,000 USD, down from $115,000 USD the previous year.

Year	% of ransoms paid	Ransomware engagements
2022-23	14	15
2023-24	13	29
2024-25	4.8	63

Year	Average ransom payment	Ransomware engagements
2022-23	130,000	15
2023-24	115,000	29
2024-25	77,000	63

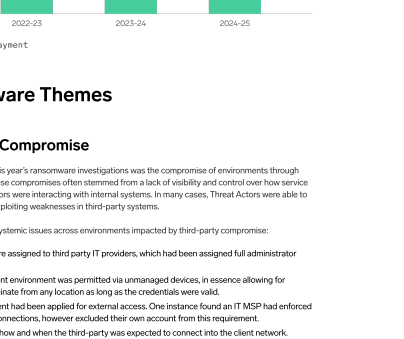

Average ransom payment

Ransomware Themes

Third Party Compromise

A recurring theme in this year's ransomware investigations was the compromise of environments through third-party access. These compromises often stemmed from a lack of visibility and control over how service providers and contractors were interacting with internal systems. In many cases, Threat Actors were able to gain initial access by exploiting weaknesses in third-party systems.

We observed several systemic issues across environments impacted by third-party compromise:

- Many accounts were assigned to third party IT providers, which had been assigned full administrator privileges.
- Access into the client environment was permitted via unmanaged devices, in essence allowing for connections to originate from any location as long as the credentials were valid.
- No MFA enforcement had been applied for external access. One instance found an IT MSP had enforced MFA for all client connections, however excluded their own account from this requirement.
- Limited visibility of how and when the third-party was expected to connect into the client network.

In many instances, external parties were being treated as trusted insiders, but without the same controls or oversight applied to internal staff.

In one case, an organisation had experienced a ransomware incident as a result of a third-party credential compromises. Following the incident, the third-party was offboarded and the project was close. Despite this, appropriate steps were still not followed to disable the access assigned to the third-party. It was subsequently found that the organisation took weeks to disable access following the cessation of the contract.

A broader contributing factor was the absence of structured identity and access management. Many organisations lacked a centralised register of third-party accounts, did not log activity associated with external users accessing internal systems, and relied on informal processes to manage onboarding and offboarding.

Lack of monitoring

One of the most consistent findings in Triskele Labs' ransomware investigations this year was the absence of action associated with responding to security alerts. In nearly every case involving unauthorised access to infrastructure, there were early indicators or alerts raised as part of the compromise. These alerts either went unnoticed, were misinterpreted, or were simply never seen.

We observed environments where Endpoint Detection and Response (EDR) tooling was deployed and functioning, yet the alerts it generated were:

- Mis-triaged or ignored by Managed Security Service Providers (MSSPs), Managed Service Providers (MSPs), or internal IT teams. In one instance, an external provider of a ransomware victim directly warned the organisation that their environment had been compromised by a ransomware Threat Actor. This warning was ignored, and the organisation was subsequently encrypted less than 24 hours later.
- Being forwarded to inboxes or ticket queues that were not being actively monitored, or were being flooded with alerts from all security and IT platforms that resulted in significant alerts being drowned out by noise.
- Generated outside of business hours, with no arrangements in place for after-hours monitoring or response.

In some incidents, the security tooling in place was not fit for purpose. Legacy antivirus or basic antimalware solutions were still being relied upon in enterprise environments, despite well-known limitations in detecting modern attacker behaviours such as credential theft, lateral movement, and scripting activity. These products were easily disabled or bypassed by Threat Actors during the early stages of compromise.

While technology was technically present in many cases, the real issue was operational. Many victims lacked defined processes, ownership, or appropriate resourcing that allowed small issues to become major breaches.

Password Hygiene

Credential misuse and poor access control remain some of the most consistent themes observed during ransomware investigations. In far too many cases, initial access or privilege escalation was gained not through sophisticated actions but through basic oversights that centred around weak passwords, shared accounts, and the absence of appropriate password controls.

It was common to encounter spreadsheets full of credentials stored in shared drives, often unprotected and accessible to anyone with network access. We regularly saw variants of the same passwords reused widely across environments, including between user and administrator accounts and across both production and backup environments. These practices made lateral movement trivial for Threat Actors once they gained an initial foothold.

Many organisations had no formal process to offboard staff or disable expired accounts, leaving unused credentials lingering for months or years. In environments without identity management, oversight of administrator accounts was often not performed at all.

The cumulative effect of these practices has led to Threat Actor leveraging basic techniques, including manually searching through shared drives, in order to elevate privileged. threat landscape where attackers don't need to be particularly sophisticated. While Triskele Labs has observed many organisations increasing their spending on cyber security controls, there are still weakness centred around the behaviour of users within environments.

Business Email Compromise (BEC) Investigations

Overview

Business Email Compromise (BEC) incidents represented 45% of all incidents handled by Triskele Labs in the 2024–25 financial year. Although the operational impact of BEC incidents is generally less severe than that of ransomware attacks, the financial consequences can be substantial, especially when payment direction fraud is involved.

These attacks responded to by Triskele Labs predominantly target victims using the Microsoft 365 (M365) suite of business productivity applications (I.E. Microsoft Outlook), reflecting its dominance as the primary platform for business communications in Australia and New Zealand.

Understanding Business Email Compromise

The vast majority of BEC incidents observed across investigations are performed by Threat Actors that are financially motivated.

- The most common outcome is payment redirection fraud, where attackers impersonate suppliers, executives, or staff to trick victims into transferring funds to attacker-controlled accounts.
- In some cases, the compromise is leveraged to scam downstream individuals or partner organisations, further amplifying the impact.

The impactful nature of financially motivated Threat Actors remained the primary driver of detection for BECs, which was most commonly detected as a result of:

- The business detecting attempted or successful payment redirection fraud conducted by the Threat Actor.
- The business is alerted to suspicious emails being sent by the Threat Actor to external parties, typically in the form of phishing emails that have been distributed in an attempt by the Threat Actor to extend their access to new victims.

BECs associated with solely information gathering objectives were rarely observed, in part due to the less impactful nature of these Threat Actors and the fact that they are less common amongst Australian small and medium businesses. Financially motivated Threat Actors continue to use techniques that result in the access and collection of information from mailbox users and the risk of this taking place as part of BECs remains, even when it is not the primary motivation of the Threat Actor.

Threat Actors are leveraging techniques designed to establish and exploit trust. Phishing campaigns typically sent once the Threat Actor has exhausted opportunities to complete payment redirection fraud within a compromised impacts, are most commonly observed issuing out phishing emails to known contacts of the compromised mailbox.

Instead of sending a malicious attachment or direction users to a suspicious URL, it is becoming increasingly common for Threat Actor to create documents within the legitimate file sharing platforms of the compromised account (I.E. SharePoint, OneDrive, Dropbox) and sharing access to the document. Receiving a link to a document hosted within a legitimate sharing location, from a known contact, is resulting in Threat Actors furthering their access to new victims.

Role of Forensics

Digital Forensics and Incident Response (DFIR) investigations are conducted with extensive experience and knowledge of the investigative process. DFIR acts as an external, trusted party focused on delivering unbiased, fact-based insights critical to understanding cyber incidents. While internal IT staff or external IT Managed Service Providers provide valuable support, often investigations do not capture the nuanced details and techniques that an investigation requires.

The objectives of a DFIR BEC investigation typically are to:

- Confirm that there had been unauthorised access to mailbox(es).
- Identify how the Threat Actor was able to gain initial access to the environment.
- Identify what actions were performed by the Threat Actor whilst they had access.
- Identify whether Personally Identifiable Information (PII) or Protected Health Information (PHI) had been accessed by the Threat Actor.
- Ensure that the Threat Actor has been purged from the environment.
- Provide guidance with the remediation and recovery activities.
- Provide recommendations to strengthen the cyber security posture post incident.

Traditional investigation methods are increasingly falling short, particularly given the growing complexity of Business Email Compromise BEC techniques. Notably, Threat Actors are now employing persistence mechanisms not observed in previous financial years, allowing them to maintain access even after standard containment and investigative measures have been applied.

Stages of a Business Email Compromise

Initial Access

Over the past financial year, Threat Actors have increasingly leveraged legitimate infrastructure, such as free hosting platforms, to host redirection and landing pages. These pages are often customised to match the branding of the compromised organisation, including their logo and business name. Because they are hosted on legitimate services and typically contain only a secondary link (rather than a direct credential harvester), they frequently bypass email filtering controls.

The delivery of phishing campaigns is widely being facilitated through legitimate services such as SharePoint, DropBox and MailChimp. Files that contain a link to a credential harvester will be uploaded to the legitimate platform, where the Threat Actor will subsequently leverage the platform's share functionality to distribute access to the file to victims. Some platforms also include the ability to schedule follow-up reminders, which can further continue the phishing emails delivery.

Threat Actors have been observed applying MFA to external SaaS solutions (I.E. DropBox, MailChimp) in order to prevent the legitimate account holders from regain access to the account.

QR code phishing (or 'quishing') remains in use, although it is currently less prevalent than the newer delivery methods described above.

Evidence Sources

As of 1 March 2025, Microsoft stopped writing audit log data directly to mailboxes. From this date forward, activity attributed to actions performed within an account is only captured within Unified Audit Log (UAL), under the condition that UAL is enabled. For Microsoft 365 tenancies created after 2021, the UAL is enabled by default. However, tenants created in 2021 or earlier will not have this enabled by default.

To check whether UAL is enabled, a privileged account holder can visit https://purview.microsoft.com and navigate to the 'Audit' section. If the UAL is not enabled, a blue banner will appear prompting you to turn it on. Enabling it may take several minutes to take effect.

Alternatively, UAL status can be verified using PowerShell:

1. Run Connect-ExchangeOnline to authenticate.
2. Enter Get-AdminAuditLogConfig to retrieve the audit log configuration.
3. Review the fields UnifiedAuditLogIngestionEnabled and UnifiedAuditLogFirstOptInDate to determine if the UAL is active and when it was enabled.
4. If not enabled, you can enable it by running: Set-AdminAuditLogConfig -UnifiedAuditLogIngestionEnabled $true

Data Exfiltration

A critical part of a Business Email Compromise (BEC) investigation is determining whether mailbox synchronisation has occurred, as this process creates a local offline copy of the mailbox. While synchronisation is designed to improve productivity by allowing users to access emails even without internet connectivity. This functionality has also introduced risk during mailbox compromises as Threat Actors can obtain a complete copy of any mailbox contents.

Delegate access is another key consideration. Within many organisations, it is common for users to be granted access to shared mailboxes or other employees' mailboxes for operational purposes. If a compromised user account has such access, the impact of the incident can quickly expand beyond a single mailbox, increasing the overall risk and complexity of the investigation.

Email clients like Microsoft Outlook automatically create an Offline Outlook Data File (OST) when a user signs in, storing emails, calendar events and contacts on the local device. This means that a Threat Actor who signs into a compromised mailbox using the Outlook desktop application can generate a full local copy of the mailbox—accessible even after containment measures are applied.

Third-party applications like eM Client provide similar functionality and are sometimes used by Threat Actors to send and receive emails from compromised accounts. eM Client supports full mailbox synchronisation, with the default configuration set to synchronise all messages excluding attachments. However, it can be adjusted to download full mailbox contents, including attachments, to the local device.

Threat Actors may also leverage legitimate cloud backup tools such as PerfectData Software to exfiltrate mailbox data. PerfectData enables complete mailbox backups from various services, including Microsoft 365 and IMAP. When used maliciously, this allows a Threat Actor to extract an entire mailbox during the window of compromise, often before any containment or recovery actions have been taken.

Persistence

One of the more advanced techniques observed in Business Email Compromise (BEC) cases for the first time this financial year was the use of Azure App Registrations to maintain persistence within Microsoft 365 environments. This technique allows Threat Actors to create an authorised application that can access mailbox data and other M365 resources via Graph API, even after credentials are reset or MFA is enforced. It is particularly effective in scenarios where traditional containment measures (I.E. password resets) are applied without a thorough investigation. In these cases, the Threat Actor can silently retain access to the environment.

In 28% of BEC cases investigated across the financial year, Triskele Labs identified some form of persistence mechanism being leveraged. Azure App Registrations were a notable trend, closely mirroring the capabilities of offensive tooling like GraphRunner, which is designed to abuse Graph API permissions. Other persistence techniques included the use of Temporary Access Passes (TAPs), the creation of new Microsoft 365 accounts, and the reactivation or password reset of dormant or inactive accounts.

These findings reinforce the need for a deeper investigative approach when responding to BEC incidents. If containment is completed without identifying and removing the mechanisms Threat Actors have established for persistence, organisations risk ongoing unauthorised access and further compromise.

Payment Redirection Fraud

Once a mailbox has been compromised, the primary motivation of the Threat Actor is typically to conduct payment redirection fraud. In many cases, the compromised organisation itself facilitates the attempt by providing the attacker with access to email conversations, invoice templates, and contact details for clients. Threat Actors use this access to impersonate the organisation in an attempt to deceive clients into transferring funds to attacker-controlled bank accounts.

However, it's increasingly common for the compromised organisation to become the victim of the fraud itself. Threat Actors are frequently registering lookalike domains to impersonate trusted suppliers or service providers, allowing them to target accounts payable teams within a legitimate business. Many organisations are drawn in through convincing emails that appear to originate from a legitimate source, with only subtle indicators of fraud. This reinforces the importance of human controls that include validating bank detail changes through a known phone number, using multi-party approval workflows for financial transactions, and training staff to identify red flags. While technical controls can reduce risk, it is often these manual validation steps that prevent fraudulent payments from being made.

Payment redirection fraud was attempted in 27% of Business Email Compromise (BEC) cases investigated by Triskele Labs over the financial year. Of those cases, Threat Actors successfully achieved their objective in 40% of attempts, highlighting the effectiveness of this tactic when detection and human validation controls fail.

On average, successful attempts resulted in the redirection of $107,927 which represented the significant financial impact that can arise from a single compromised mailbox or supplier impersonation.

Detection

Most BEC incidents are first identified in one of two ways:

- Receipt of an external notification following phishing emails sent from the compromised account; or
- An internal discovery due to delayed payments or financial discrepancies.

The delayed and reactive nature of these detection methods means Threat Actors are often able to operate within the environment unnoticed for an extended period. Across cases investigated by Triskele Labs, the average dwell time before detection was 32 days (down from 41 days in the previous financial year), with most incidents only identified once the attacker took an action that caused visible impact.

Payment Redirection Fraud

Once a mailbox has been compromised, the primary motivation of the Threat Actor is typically to conduct payment redirection fraud. In many cases, the compromised organisation itself facilitates the attempt by providing the attacker with access to email conversations, invoice templates, and contact details for clients. Threat Actors use this access to impersonate the organisation in an attempt to deceive clients into transferring funds to attacker-controlled bank accounts.

However, it's increasingly common for the compromised organisation to become the victim of the fraud itself. Threat Actors are frequently registering lookalike domains to impersonate trusted suppliers or service providers, allowing them to target accounts payable teams within a legitimate business. Many organisations are drawn in through convincing emails that appear to originate from a legitimate source, with only subtle indicators of fraud. This reinforces the importance of human controls that include validating bank detail changes through a known phone number, using multi-party approval workflows for financial transactions, and training staff to identify red flags. While technical controls can reduce risk, it is often these manual validation steps that prevent fraudulent payments from being made.

> Payment redirection fraud was attempted in 27% of Business Email Compromise (BEC) cases investigated by Triskele Labs over the financial year. Of those cases, Threat Actors successfully achieved their objective in 40% of attempts, highlighting the effectiveness of this tactic when detection and human validation controls fail. On average, successful attempts resulted in the redirection of $107,927 which represented the significant financial impact that can arise from a single compromised mailbox or supplier impersonation.

Detection

MostBEC incidents are first identified in one of two ways:

- Receipt of an external notification following phishing emails sent from the compromised account; or
- An internal discovery due to delayed payments or financial discrepancies.

The delayed and reactive nature of these detection methods means Threat Actors are often able to operate within the environment unnoticed for an extended period.Across cases investigated by Triskele Labs, the average dwell time before detection was 32 days (down from 41 days in the previous financial year), with most incidents only identified once the attacker took an action that caused visible impact.

Payment Redirection Fraud

Once a mailbox has been compromised, the primary motivation of the Threat Actor is typically to conduct payment redirection fraud. In many cases, the compromised organisation itself facilitates the attempt by providing the attacker with access to email conversations, invoice templates, and contact details for clients. Threat Actors use this access to impersonate the organisation in an attempt to deceive clients into transferring funds to attacker-controlled bank accounts.

However, it's increasingly common for the compromised organisation to become the victim of the fraud itself. Threat Actors are frequently registering lookalike domains to impersonate trusted suppliers or service providers, allowing them to target accounts payable teams within a legitimate business. Many organisations are drawn in through convincing emails that appear to originate from a legitimate source, with only subtle indicators of fraud. This reinforces the importance of human controls that include validating bank detail changes through a known phone number, using multi-party approval workflows for financial transactions, and training staff to identify red flags.While technical controls can reduce risk, it is often these manual validation steps that prevent fraudulent payments from being made.

Payment redirection fraud was attempted in 27% of Business Email Compromise (BEC) cases investigated by Triskele Labs over the financial year. Of those cases, ThreatActors successfully achieved their objective in 40% of attempts, highlighting the effectiveness of this tactic when detection and human validation controls fail.

On average, successful attempts resulted in the redirection of $107,927 which represented the significant financial impact that can arise from a single compromised mailbox or supplier impersonation.

BEC Themes

Session Token Theft

While the uptake of Multifactor Authentication (MFA) continues slowly across Australia, Across Australia and New Zealand, many businesses are implementing MFA as there it has been a common belief that MFA will mitigate the risk of credential compromise. Unfortunately, Threat Actors are continuing to embed techniques within their technology to bypass these MFA controls, particularly within the Microsoft 365 platform due to it's prevalence of use across the industry.

Session token theft is a form of Adversary-in-the-Middle (AitM) attack whereby a Threat Actor will develop a credential harvester that impersonates the Microsoft 365 authentication portal. As a victim enters their legitimate username, password and MFA token within the credential harvester, the information is used to legitimate authenticate into Microsoft 365 in order to collect an authorised session token.

Many users are on the lookout for suspicious emails or poorly coded authentication portals, however the technologies in use by Threat Actors are bypassing these perceptions of the public and resulting in successful email compromises.

In one instance across FY25, the leader of a small business has been championing cyber awareness within their organisation and led a 12 month education and training program amongst staff. At the end of this program, the leader fell victim to a phishing campaign that leveraged common techniques. Information provided by the leader found that these common techniques had not been included as part of the provided training.

The techniques used by Threat Actors are streamlined and being used frequently to conduct attacks which can be shared with staff easily to increase awareness and build the cyber resilience of organisations.

Mobile Devices

While the convenience offered by accessing emails from your phone is seen as a positive for many business leaders, it also introduces an additional risk when it comes of the proliferation of BECs and payment redirection fraud.

- Many common email applications designed for mobile devices will only show the senders email name, not the full email address. This has been a contributing factor in many BECs across FY25, where the recipient of the email does not identify that a Threat Actor is impersonating the email address of an external contact.
- It is common for staff to use personal mobile devices which are not managed by the organisation.
- Actions associated with interactions with phishing emails or malicious attachments are not detected by the organisation's security tooling.
- Some Threat Actor continue to leverage QR codes as a mechanism to deliver credential harvesters, which are designed to initiate activity from mobile devices, instead of managed laptops and endpoints.

Excessive Permissions

Particularly common across small businesses, Triskele Labs has found that many businesses and user accounts are being assigned excessive (or top level) administrative privileges within email environments.

Larger organisations have mature processes in place that necessitate the implementation of Identity and Access Management, Principles of Least Privilege and dedicated Service Accounts for administrative users. For smaller businesses, it is not uncommon for leaders or offices administrators to be assigned administrative permissions that are commonly being exploited by Threat Actors.

In one instance recently responded to by Triskele Labs, an IT user was assigned temporary top level privileges within their Microsoft 365 tenancy in order to complete an upgrade project. During these period of elevated access, the user's account was compromised by a Threat Actor. Leveraging the top-level privileges, the Threat Actor commandeered the entire Microsoft 365 and Azure Cloud tenancy of the business and locked legitimate business users out of the environment entirely.

Many users assigned top-level privileges (I.E. Microsoft 365's Global Administrator role) do not require the full use to the assigned permissions and would benefit from being assigned a role with fewer account privileges.

Even for standard user accounts, there are privileges commonly assigned that are being exploited by Threat Actors to enable the rapid access and collection of information from the mailbox user. Microsoft 365 allows for users to authenticate to third-party applications in order to interact with their accounts in different ways; Threat Actors are frequently exploiting this permission and authenticating to applications designed to collect copies of mailboxes, extract information from mailboxes or issue mass phishing campaigns from mailboxes.

Key BEC Controls

There are some common controls Triskele Labs recommends to mitigate the likelihood and impact of BECs affecting organisations:

- Enforce Multi-Factor Authentication to protect again unsophisticated attack methods (I.E. credential stuffing).
- Implement policies to restrict access to resources from outside of expected locations or services (I.E. Geoblocking policies, third-party applications, unauthorised devices)
- Segregate administrative access to platforms to dedicated non-user accounts.
- Assess the opportunities to implement session token protection controls.
- Implement 24x7 monitoring in order to detect unusual sign-in activities.
- Govern how third-party administrative access is being assigned and monitored.

Emerging Risks

Lack of Encryption

The emergence of ransomware Threat Actors that don't encrypt data. Throughout FY25, Triskele Labs identified two (2) Threat Actors that obtained sufficient access to systems to allow for encryption to be deployed, however the Threat Actor was observed actively choosing not to encrypt victim environments:

- One Threat Actor made assertions of their leak site that while they did compromise and extort victims across all industry verticals, they did not encrypt the environments of healthcare providing organisations. Triskele Labs noted that this Threat Actor followed this rule on one (1) occasions, whilst breaking in earlier in the year for other healthcare providing victims.
- A new Threat Group emerged in FY25 that does not seek to encrypt systems. Instead of encrypting systems and making them unavailable, this Threat Actor solely exfiltrates data and relies on organisations seeking to pay in order to prevent the data from being published.

Further to this, examples were identified of Threat Actors, notably KillSec3, conducting exfiltration attacks against publicly accessible data stores and vulnerable File Transfer Protocol (FTP) servers. During these attacks, the Threat Actor sought to exfiltrate unprotected data from these locations and publish the data on their Dark Web leak site. In manage instances the organisation was unaware that the incident had occurred until the name of the organisation was listed on a Dark Web leak site.

With many organisations not applying appropriate controls to personal, sensitive or health information, Threat Actors are presented with less barriers to exfiltrate data than they are to encrypt systems.

Key Controls:

- Network restrictions to prevent open internet access from infrastructure.
- Adoption of Data Loss Prevention (DLP) controls.
- Strong Identify and Access Management (IAM) controls restricting unnecessary access. permissions within the environment.
- Data encryption at risk.
- Secure configuration and management of file storage and file transfer infrastructure.

Endpoint Compromise

Throughout FY25, Triskele Labs responded to multiple cyber incidents that originated from endpoint compromises.

Key Controls

Unreported Red Incidents

Key Controls

Conclusion

GRC

If you only look at the flashy headlines, the cyber security industry is winning. Ransomware numbers are down. RedLine stealer has been dismantled. The head of LockBit identified. The Medibank hacker (sic) publicly shamed. But in an industry where winning and losing is defined by one question - are we enabling all companies to become more secure, the reality is far more complicated.

The Illusion of Progress

Why Australia's businesses are still losing the cyber war

If you only look at the flashy headlines, the cyber security industry is winning. Ransomware numbers are down. RedLine stealer has been dismantled. The head of LockBit identified. The Medibank hacker (sic) publicly shamed. But in an industry where winning and losing is defined by one question - are we enabling all companies to become more secure, the reality is far more complicated.

The growing divide

The divide between companies with effective and ineffective cyber security programs is significant. The top players have the expertise and the resources to engage with Government agencies and contribute to the profession.

The rookies aren't even making it past the starting line. Snowballing compliance burdens and ill-fitting expert advice are preventing companies with fewer resources from making progress. Much to the exasperation of the Federal Government, small- and medium-sized businesses remain naïve and inexperienced regarding the dangers they face using modern I.T. infrastructure.

> "The divide between companies with effective and ineffective cyber security programs is significant."

Lack of Awareness

The family-run winery does not understand the significance of allowing lax configurations to remain in their environment. What does "disabling legacy authentication in Microsoft365" mean anyway?

Adding to this lack of expertise, over the last 12 months we have witnessed the troubling trend of "Trickle Down Cyber Security." As larger companies become more regulated, they are implementing third-party risk management programs and pushing onerous compliance burdens down the supply chain.

Compliance without context

Does your small- or medium-sized business want to bid on a contract? Please send over your ISO 27001 certificate, SOC 2 Type 2 Report, and your business continuity plan. This is a common introduction to cyber security for many small- or medium-sized Australian businesses if they have been lucky enough to avoid getting hacked.

Now aware of the requirements, the person at the SMB who was unfortunate enough to Google ISO 27001 is tasked with implementing the ISMS, organising certification, and responding to the UpGuard questionnaires.

The net result is a cumbersome and unworkable information security framework implementation. Significant resources spent to win a contract, but lack of attention to basic technical controls means the business is no more secure than it was before the addition of 30 new policies and procedures.

> "Significant resources spent to win a contract, but lack of attention to basic technical controls means the business is no more secure than it was before."

Industry accountability

The cyber security industry is also to blame for this boondoggle. Too many organisations are out looking to make a quick buck and provide generic guidance without taking the time to understand the organization's I.T. infrastructure or how their business operates.

Governments must take a hard look at the programs and initiatives they are rolling out. If the intended recipients don't understand how to apply the advice, or worse yet, even where it find it, it's a wasted effort. Guidance needs to be written in plain English, contextualised to the size and maturity of the business, and delivered through channels that business owners actually use.

Threat actors are aware of this security gap, and they are looking for low-hanging fruit. They need their payday, after all. When threat actors discover that no multi-factor authentication exists on remote connection services, firewall access is exposed on the internet, or systems are out-of-date and insecure, they see dollar signs. Smaller companies make easier targets, and they tend to suffer more severe consequences.

Lifting the supply chain

If we want supply chain security to meaningfully improve, larger organisations must invest in educating their suppliers, providing funding or resources where possible, and adjusting their expectations based on the realities of operating a small business.

Otherwise, we're just creating compliance theatre that checks a box but doesn't lift the security needle. True uplift means rolling up sleeves and getting involved, and not outsourcing accountability to the nearest procurement team. Just the other day the Victorian Government was asking a three-person startup if they had completed an IRAP assessment.

Lastly, the security community must do better in making cyber advice accessible and relevant. That starts by understanding that a five-person accounting firm doesn't have a CISO and can't afford a GRC consultant. Advice needs to be pragmatic, actionable and based on real-world environments, not ideal-state architectures or what the ISO standard says.

The real challenge

If we can't bridge the gap between what's technically perfect and what's practically possible, we'll keep seeing the same headline: "Family business crippled by ransomware, pays ransom".

This is the biggest challenge the industry faces now.

Provide practical, easy to implement cyber security programs to those businesses that do not know they need them.

> "Smaller companies make easier targets, and they tend to suffer more severe consequences."

MDR

A year of smarter detection, faster response and measurable improvement.

SOC Performance and Growth
Record alert volumes, expanding client base, and continued operational efficiency through automation and tuning.
Client Growth and Sector Expansion
Efficiency at Scale
Context in the National Threat Landscape
Detection and Response Performance
Rapid detection, targeted response, and continued refinement of alert quality underpin SOC efficiency and resilience.
Detection Efficacy
Alert Quality and False Positives
Mean Time to Detect (MTTD)
Mean Time to Respond (MTTR)
Mean Time to Conclusion (MTTC)
High MTTC
Cases by Source
SIEM remains the foundation of SOC

SOC Performance and Growth

Record alert volumes, expanding client base, and continued operational efficiency through automation and tuning.

Client Growth and Sector Expansion

In FY25, the Triskele Labs Security Operations Centre (SOC) processed approximately 263,000 alerts — a significant increase from nearly 185,000 in FY24. This growth aligns with a 69% year-on-year rise in active Managed Detection and Response (MDR) clients across fully managed and DFIR monitoring services.

The SOC's reach expanded across several new sectors, including construction, education and logistics, while maintaining strong representation in established industries such as finance and insurance, health and community care, professional and business services, retail, FMCG and consumer services, not-for-profit, and industrial sectors.

Efficiency at Scale

Despite a sharp increase in total alerts, the average volume per client remained steady at around 222 alerts per month. This stability underscores the effectiveness of Triskele Labs' ongoing investments in correlation tuning, enrichment pipelines, and suppression logic.

When isolating manually triaged cases, excluding those managed by in-house automation, the figure reduced further to an average of 181 alerts per client per month, reflecting the success of automation initiatives in improving efficiency and analyst focus.

Detection and Response Performance

visibility, enabling comprehensive threat detection beyond the endpoint.

Overview

Case priority breakdown

Strategic Importance of SIEM

Complementary Role of EDR

Interesting Statistics

Alert volumes continue to grow while detection accuracy and true positive rates improve year-on-year.

Overview of Alert Volumes

Ticketing and Escalation Trends

Determination Accuracy and Quality

Key Takeaways

Top 10 Cases

Initial access and authentication anomalies remain dominant vectors inclient environments.

Overview

Strategic Insight

Cases by Vertical

Finance, Health, and Critical Infrastructure sectors remain the most targeted industries under SOC monitoring.

Overview

Strategic Insight

Rapid detection, targeted response, and continued refinement of alert quality underpin SOC efficiency and resilience.

Detection Efficacy

The Triskele Labs SOC continues to demonstrate strong internal detection capability, identifying and acting on critical incidents within a median of 7.3 minutes. This performance reflects significant investment in tooling, automation, and analyst workflows that enable rapid visibility and containment across client environments.

For context, Mandiant's M-Trends 2023 report indicates that only 37% of breaches are typically detected internally across most industries, with the majority identified by external parties. Triskele Labs' consistently low Mean Time to Detect (MTTD) highlights the maturity of its detection engineering and triage models, particularly across higher-risk sectors such as finance, healthcare, and professional services.

Alert Quality and False Positives

Alert fidelity remains a strategic focus. Internal analysis indicates that approximately 70% of alerts processed do not translate into actionable security incidents, a proportion consistent with broader industry experience.

Through targeted rule tuning, suppression logic, and contextual enrichment, the SOC continues to reduce noise and direct analyst attention toward meaningful alerts. Client-specific suppression rules and enhanced routing within the case management platform have further improved accuracy and workflow efficiency.

These measures have maintained analyst throughput and helped mitigate operational fatigue, even as telemetry volumes increase. Continuous tuning remains central to sustaining detection accuracy and response speed.

Mean Time to Detect (MTTD)

Median MTTD for critical incidents improved to **7.3 minutes** in FY24–25, down from 9.1 minutes in FY22–23. High-priority incidents averaged between **18–20 minutes**, remaining well within benchmarked industry performance.

This improvement reflects Triskele Labs' ongoing investment in detection engineering, threat model refinement, and enriched telemetry pipelines that accelerate context delivery to analysts.

Mean Time to Respond (MTTR)

Median MTTR for critical incidents fell to **3.6 minutes**, improving on **4.7 minutes** the previous year. This metric captures the period between detection and containment action, such as host isolation or user disablement, reflecting the speed and precision of containment workflows.

For high-priority cases, MTTR rose modestly to **5.6 minutes** from **3.5 minutes**. This change is attributed to triage sequencing, with critical incidents prioritised first and high cases requiring deeper investigation. It also reflects success in tuning and enrichment, where non-actionable alerts are increasingly filtered before analyst engagement.

Mean Time to Conclusion (MTTC)

MTTC, the sum of detection and response time, remained well within internal thresholds of **≤60 minutes for critical** and **≤120 minutes for high** incidents.

FY24: 7.29 +3.60 = 10.89 minutes
FY25: 7.27 +3.72 = 10.99 minutes

High MTTC

FY24: 19.73 +3.56 = 23.29 minutes
FY25: 19.38 +5.62 = 25.00 minutes

With internal MTTC thresholds of ≤60 minutes for critical and ≤120 minutes for high, Triskele's SOC remains well under target across both categories.

These figures demonstrate consistent containment speed and reinforce the operational maturity of Triskele Labs' SOC model — balancing automation, detection precision, and human response.

Cases by Source

SIEM remains the foundation of SOC visibility, enabling comprehensive threat detection beyond the endpoint.

Overview

Analysis of FY25 case distribution shows that **87%** of alerts originated from the **Security Information and Event Management (SIEM)** platform, **12%** from Endpoint Detection and Response (EDR)tools, and 1% from Dark Web monitoring. Phishing triage contributed a negligible volume during the period.

These figures reinforce that while EDR solutions play an important role in endpoint-level visibility, SIEM remains the backbone of effective Security Operations Centre (SOC) and Managed Detection and Response (MDR) capabilities.

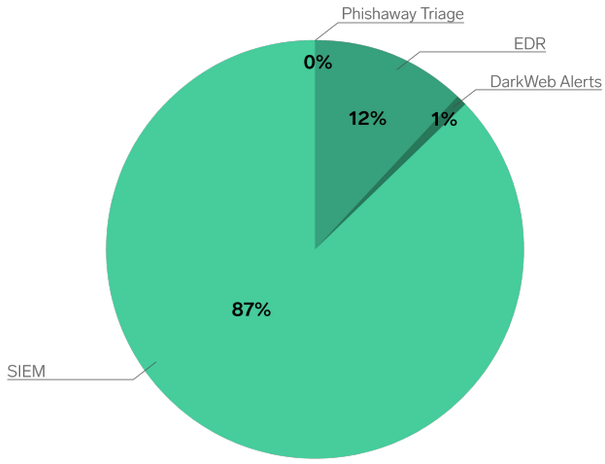

Case source distribution showing SIEM as the primary detection engine, providing comprehensive visibility beyond endpoint telemetry.

Case priority breakdown

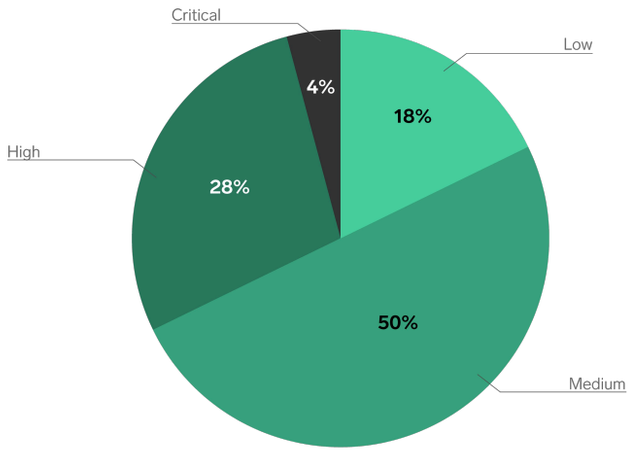

Non-actionable alerts represent expected activity within monitored environments — operationally relevant but with minimal security impact. As shown below, nearly 75% of cases were prioritised by their originating tools as Low or Medium severity.

Strategic Importance of SIEM

Modern cyber threats rarely remain confined to a single endpoint. Attackers increasingly exploit credentials, move laterally through networks, and target cloud and SaaS infrastructure. In this landscape, an endpoint-only view is insufficient.

A well-tuned SIEM aggregates and correlates telemetry from across the environment, including identity providers, network devices, servers, cloud applications, and security controls. This unified perspective enables analysts to detect cross-domain activity and identify high-impact threats that would otherwise go unnoticed in isolated systems.

Complementary Role of EDR

EDR platforms remain an essential component of Triskele Labs' detection ecosystem, providing deep forensic insight and rapid containment capability at the device level. However, their scope is inherently limited. Without correlation against broader environmental data through the SIEM, incidents that span multiple vectors, such as lateral movement or privilege escalation, may evade detection.

The strong ratio of SIEM-driven detections highlights the continued value of a correlated, context-aware defence model that integrates endpoint, network, and cloud telemetry into a single operational picture.

Interesting Statistics

Alert volumes continue to grow while detection accuracy and true positive rates improve year-on-year.

Overview of Alert Volumes

Total alerts processed by the SOC increased slightly to 247,922 in FY25, maintaining stability following a major uplift in FY24. The majority of cases were categorised as Medium and **Low** severity, reflecting a mature and proactive monitoring posture rather than a reactive one.

High-severity alerts accounted for approximately **16%** of total cases, while **Critical** alerts continued to represent a small but strategically important proportion at around **2%**. This profile aligns with expected trends in mature SOC environments, where strong preventive controls reduce the frequency of high-impact incidents.

Ticketing and Escalation Trends

The total number of ticketed cases rose modestly to **10,442,** representing **4.2%** of total alerts. While this percentage is lower than FY24's 5.6%, it reflects improvements in automation, correlation, and suppression logic that are successfully filtering non-actionable alerts before analyst intervention.

Ticketing rates for **Critical** incidents remained high at 16%, consistent with operational prioritisation policies, while **Medium** and **Low** cases saw a proportional decline due to enhanced rule tuning and enrichment processes. This indicates that the SOC is increasingly focused on high-value, actionable events rather than volume-driven escalations.

Determination Accuracy and Quality

Accuracy metrics show continued improvement in detection quality. The proportion of **True Positive** determinations rose to **40.08%**, up from **34.35%** in FY24 and 17.84% two years prior, a clear indicator of ongoing refinement in correlation rules and enrichment accuracy.

Meanwhile, **False Positive** rates decreased to 59.52%, continuing a downward trajectory year-on-year. This reduction highlights the effectiveness of continuous tuning initiatives and the integration of automation in early-stage triage. Indeterminate outcomes dropped to just 0.39%, underscoring improved clarity in alert disposition and case resolution.

Key Takeaways

- **Stabilised alert growth** following last year's expansion suggests sustained operational maturity.
- **Improved detection accuracy** reflects investment in tuning, enrichment, and automation.
- **Reduced false positives** have directly enhanced analyst efficiency and reduced fatigue.
- **Critical and high alerts remain tightly managed,** demonstrating effective prioritisation and triage discipline.

Top 10 Cases

Initial access and authentication anomalies remain dominant vectors inclient environments.

Overview

Analysis of the top 10 case types across all clients shows that initial access and credential-based anomalies continue to represent the most significant threat categories. The two leading detections — A*nomalous Network Behaviour (PortAccess) and Multiple Country IngressAuthentications*, together accounted for nearly **half of all observed cases.**

These findings reinforce that adversaries are frequently exploiting authentication pathways and misconfigured access controls to establish footholds within client networks. The prominence of "First Ingress Authentication from Country" and "Attacker Behaviour Detected" further highlights the sophistication of lateral movement and reconnaissance techniques observed throughout the year.

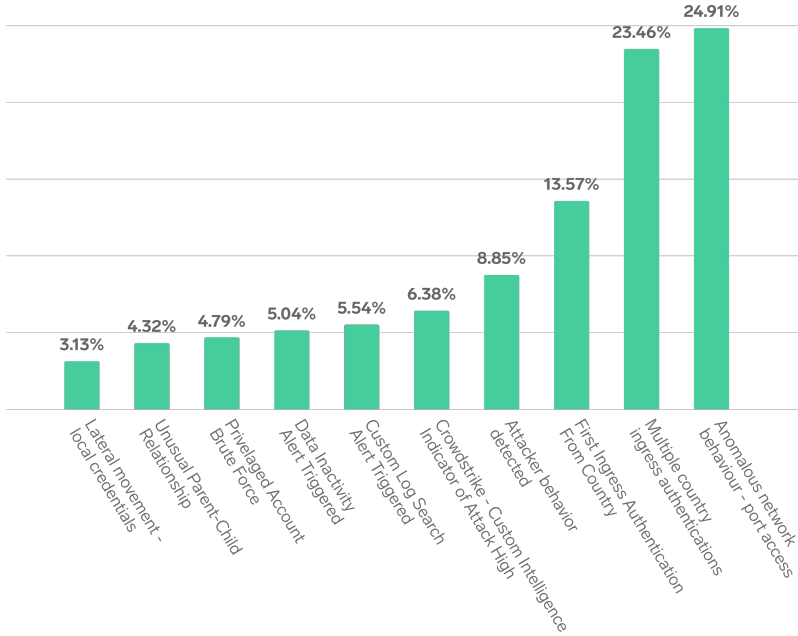

Top 10 detections across clients highlight initial access, credential misuse, and anomalous network behaviour as key risk drivers in FY25.

Strategic Insight

This pattern underscores that while endpoint and network defences remain effective at identifying known signatures and behavioural anomalies, identity-focused attacks are now a primary risk vector. Stronger control frameworks around multi-factor authentication, geolocation policies, and privileged account monitoring continue to be essential in mitigating this risk.

The persistence of *Lateral Movement — Local Credentials* and Brute Force Activity within the top ten reinforces the need for ongoing investment in **detection engineering and identity telemetry correlation** across clients.

Cases by Vertical

Finance, Health, and Critical Infrastructure sectors remain the most targeted industries under SOC monitoring.

Overview

Case distribution by industry shows the **Finance and Insurance** sector continues to record the highest case volume, representing approximately **20%** of all monitored activity. This is followed by Health and Community Care (11.7%), Technology and Engineering (10.6%), and Professional and Business Services (9.1%).

These sectors consistently appear among the most targeted across Australia due to their **data sensitivity, operational criticality, and interconnected digital ecosystems**. Finance and Health in particular face sustained targeting from both financially motivated and state-aligned actors.

The **Industrial and Resources** and **Consumer Services** sectors also contributed notable case volumes, reflecting growing attacker interest in operational technology (OT) and supply-chain environments.

Strategic Insight

This distribution aligns with national threat intelligence trends, where adversaries increasingly prioritise sectors that combine high-value data with complex, decentralised IT and OT environments. Continuous monitoring, sector-specific detection content, and adaptive enrichment logic remain vital in addressing this evolving exposure.

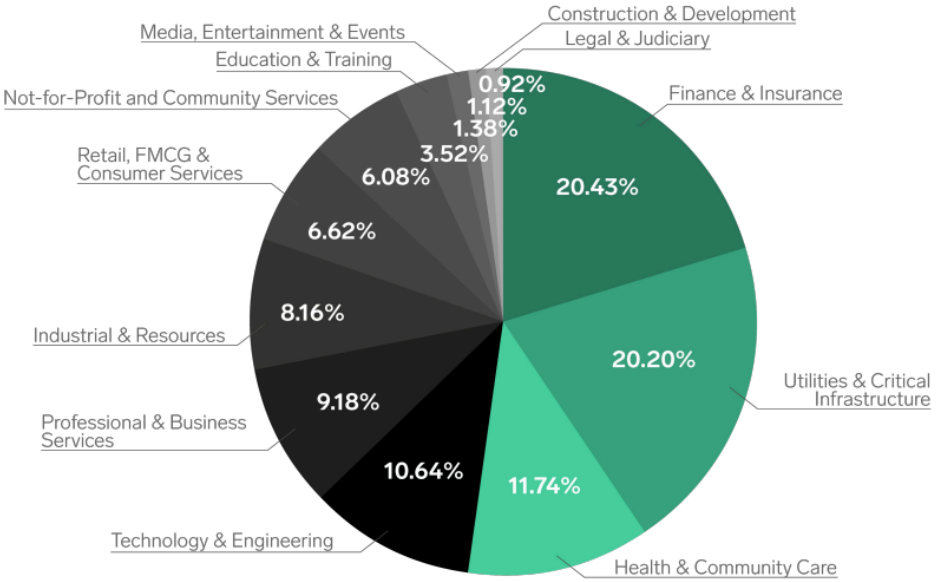

Case distribution by vertical showing financial, health, and industrial sectors as the most frequently targeted under Triskele Labs' monitoring in FY25.

© All Rights Reserved. Triskele Labs

Offensive

Triskele Labs have the privilege of working with clients through a multitude of sectors, including Government & Defense, Financial Services, Industrial, Healthcare & Life Sciences, Retail & E-commerce, Information Technology & Telecommunications, Energy & Utilities, and Education & Research.

Introduction

Eight vulnerabilities per engagement

Triskele Labs have the privilege of working with clients through a multitude of sectors, including Government & Defense, Financial Services, Industrial, Healthcare & Life Sciences, Retail & E-commerce, Information Technology & Telecommunications, Energy & Utilities, and Education & Research.

During the last financial year, as global cyber threats intensified, our team conducted over 500 engagements. Although the number and severity of findings varied considerably across these assessments, we identified a total of 3,887 vulnerabilities, an average of nearly eight per engagement and fifteen new vulnerabilities each day.

Upon further analysis, several notable trends emerged. The incidence of critical vulnerabilities has declined, while the number of longstanding clients engaged in continuous testing has grown. This relationship highlights the direct impact of regular penetration testing and the implementation of its recommendations on reducing overall risk within client environments.

Severity analysis

All identified issues are assigned a risk rating to allow for prioritisation of remediation efforts. Triskele Labs utilises a custom risk rating framework to identify risk level, built from tried and trusted processes and experience. The Triskele Labs framework is detailed below.

	Impact		
	Major	Moderate	Minor
Likely	Critical	High	Medium
Possible	High	Medium	Low
Unlikely	Medium	Low	Low

These Can be derived from the following principles:

Rating	Example Characteristics	Resolution
Critical	• Previous compromise • Definite sensitive information disclosure • Critical systems impacted • Complete compromise • Trivial exploitation • Authentication bypasses	Within 7 days
High	• Possible sensitive information disclosure • User account compromise • Trivial exploitation • No countermeasures	Within 30 days
Medium	• Possible limited information disclosure • Requires a skilled attacker • Custom exploits or code • Some countermeasures in place	Within 90 days
Low	• Critical data not at risk • Requires several attack vectors • Difficult to exploit • Countermeasures in place to prevent	Within 180 days
Info	• Notification of a quirk or bug with no impact • Best practice recommendation that may improve security posture.	

- **Risk:** Likelihood × Impact.
- **Likelihood:** Vulnerability × Threat.
- **Vulnerability:** An error or weakness in the design, system implementation or system.
- **Threat:** A person (internal/external) or something that can cause damage to an asset.

Impact: Loss of confidentiality, integrity, availability & accountability. It can consist of ease of vulnerability being detected, exploited, and discovered or information leakage.

Likelihood definitions

UNLIKELY: Although a flaw or weakness exists, this flaw is either exceedingly difficult to exploit, would require a highly skilled attacker, needs to be combined with other vulnerabilities in order for exploitation to occur, or is otherwise affected by factors which reduce the likelihood, such as timing or conditions.

POSSIBLE: An attacker with a moderate level of skills and knowledge may be able to exploit this flaw or weakness, however the flaw is not widely exploited.

Likely: The flaw is well-know and exploitable by a novice or inexperienced attacker.

Consequence definitions

Minor: The result of exploiting the flaw or weakness is a small gain for the attacker. This could include access to information that may assist other attacks which is not highly sensitive. The business should experience little impact as a result of exploitation of this flaw.

Moderate: Exploitation of the flaw would result in the attacker gaining access a low level of access to a resource, or access to sensitive information, however there are limitations and restrictions on this access.

Major: Exploitation of the flaw results in a high level of access to the attack, or a major change impacting the business. This is a position from which significant damage can be caused. This includes access to highly sensitive information.

Vulnerability analysis

The following provides a broad overview of trends and assessments conducted by Triskele Labs.

Metric	Prior to 2024	Post 2024
Average vulnerabilities per assessment	7.5	7.76
Assessments with at least 1 Critical vulnerability	11	6
Assessments with at least 1 High vulnerability	82	138

Severity Breakdown	Prior to 2024	Post 2024
Critical	18 (0.8%)	10 (0.3%)
High	198 (8.3%)	304 (8%)
Medium	573 (24.1%)	872 (22.8%)
Low	1219 (51.4%)	2072 (54.2%)
Informational	365 (15.4%)	567 (14.8%)

Key takeaways:

- Overall volume and density are almost identical - 7.5 to 7.76 vulns per workspace in both timeframes.
- Critical issues are proportionally rare post-2024 (0.3 % vs 0.8 %), and fewer engagements identified any Critical findings.
- High severity issues remained at ~8 % of all vulnerabilities for both periods, however a higher number of engagements post 2024 (138 vs 82) saw at least one High finding, reflecting the larger sample.
- Medium & Informational vulnerabilities reduced slightly in share, while Low risk findings rose from ~51 % pre-2024 to ~54 % post-2024.

Most Common Test Types

Consistent with previous years, the three most common types of penetration testing include the following:

Web Applications and Application Programming Interfaces (APIs)

Designed to identify and exploit vulnerabilities in web applications and their supporting APIs, from small, single page brochure style web applications, to large, complex, Customer Relation Management (CRM) software. The primary goal is to evaluate the

application's resilience against common threats. This testing involves a combination of automated scanning and manual techniques to assess business logic, input validation, access controls, and more. Inclusions typically cover OWASP Top 10 vulnerabilities.

External Networks

Focuses on assessing the security posture of an organization's internet-facing assets, such as firewalls, VPN gateways, email servers, and public IP ranges. The primary goal is to identify and exploit vulnerabilities that an external attacker could leverage to gain unauthorized access to the internal network or sensitive systems.

Internal Networks

A controlled assessment that simulates an attacker who has already gained access to an organization's internal environment, whether through physical access, social engineering, or as a rogue employee, in what is known as an "assumed Breach" scenario. The primary goal is to evaluate the effectiveness of internal security controls in preventing privilege escalation, lateral movement, and unauthorized access to critical systems and sensitive data

OWASP Top 10

Using the OWASP Top-10 is perhaps the most effective first step towards changing the software development culture within organizations into one that produces more secure outcomes.

Our 2024 analysis highlights five security gaps as repeat offenders:

- Broken Object Level Authorisation has risen to the top as attackers still exploit simple ID manipulation in APIs and web applications to harvest or tamper with data, placing customer records at constant risk.
- Broken Authentication continues to enable credential stuffing and session hijacking, turning weak login flows and long-lived service tokens into fraud and compliance penalties
- Server-Side Request Forgery remains critical because cloud workloads that fetch URLs on behalf of users allow adversaries to reach internal metadata services and steal keys that can compromise entire environments
- Security Misconfiguration still plagues mature teams, with default credentials, open buckets, and verbose error pages giving intruders easy footholds that quickly escalate to full outages
- Improper Assets Management exposes untracked or outdated hosts, leaving known vulnerabilities unpatched and hidden data flows unmonitored so lack of real-time asset inventory can undermine every other control.

Emerging Sectors

Two key emerging fields which are experiencing the highest rate of testing growth are environments which contain Operational Technology (OT).

Transport, healthcare, and infrastructure are growing at a rapid pace, and are naturally a more prominent target for cyber security related threats. As these sectors routinely utilising unique hardware for highly specialised tasks and capabilities, conducting frequent and thorough testing has become paramount to ensuring operational stability, and protection of customer information.

While Operational technology holds significant importance, not all specialised tools are "air-gapped" or sufficiently segmented. This has a direct relationship with the second emerging field, being red team, or "adversarial simulation" style testing. This testing ensures that, alongside all the default internal network, active directory, and network segmentation best practices are well designed, architected, and adhered to, but that the Security Operations Centre teams are operating with full visibility and responsiveness to incoming threats of a more stealthy and targeted approach.

These tests have highlighted the importance of adherence to the more baseline principles such as credential management, Multi-Factor Authentication (MFA), which are still the most common access vector in the majority of real world attacks, and penetration testing exercises.

The Emergence of AI In Cyber Security

An ever-increasing aspect within cyber security is the implementation of Artificial Intelligence, Machine Learning, and Large Language Models within the security landscape. This pertains not only to the creation of applications via LLM assisted codebase creation, but also within defensive and offensive tools used to conduct automated testing and automated discovery. These tools should be treated with caution, as they are disruptive and may introduce a multitude of potential vulnerabilities, with the emphasis on creating a "working baseline" with no emphasis on security best practices.

About:

The Offensive team at Triskele Labs is recognised as an industry leader in identifying and sharing critical vulnerability insights on both local and global scales. Supported by highly qualified specialists, we conduct in-depth trend analyses to ensure comprehensive coverage of both established and emerging techniques within the continually evolving threat landscape.

Our expert penetration testers work diligently to ensure the efficacy of security implementations of clients' networks and environments, both physical and virtual. The main priority of penetration testing is to reveal potential entry points before malicious actors can exploit them and provide our clients with clear, prioritized, and actionable recommendations for remediation.