Commissioner Carly Kind
Australian Privacy Commissioner
Office of the Australian Information Commissioner
GPO Box 5288
Sydney NSW 2001
AUSTRALIA

10th of August 2026

Threat Intelligence and the Qantas Breach

Commissioner Carly Kind,

I write about one part of the OAIC's reasoning in its report on the 2025 Qantas data breach.

The preliminary inquiries revealed no omissions or failings warranting a Commissioner-initiated investigation (CII). The OAIC relied partly on the view that Qantas could not reasonably have foreseen and prevented the breach as it occurred.

However, the report does not address important counterfactuals, including whether public threat intelligence should have alerted Qantas to the attack and enabled targeted mitigation.

Furthermore, although preliminary and not an endorsement of Qantas, the report's reasoning could give organisations facing future scrutiny a ready defence: that the attacker used an uncommon tactic or their software offered no control capable of preventing the breach.

To assess this fairly, I first reconstruct the OAIC's argument. I then explain where the published reasoning falls short, the risk this creates for future cases, and an alternative approach.

You considered whether Qantas could reasonably have foreseen and prevented the breach:

> *"Based on the information provided, it does not appear that Qantas could have reasonably foreseen and prevented the breach in the manner that it occurred."*

Based on the preliminary inquiries, you reported that the less common tactic of inducing an agent to authorise a third-party application was not indicative of a systemic staff-training deficiency, and that Qantas's RBAC would not have prevented the breach:

> *"The way in which the threat actor gained access was through a vishing attack which could not have been prevented by a strengthening of Qantas' current role based access controls."*

Taken at its strongest, this part of your reasoning was that Qantas had controls consistent with recognised practice, while the attack used a tactic that standard training and stronger RBAC would not ordinarily have addressed:

1.1. Standard social-engineering training generally focuses on credential theft rather than inducing legitimate system actions.

1.2. Qantas had recurring training, audits, access controls, monitoring and incident-response arrangements.

1.3. Strengthening Qantas's RBAC would not have prevented the agent from authorising the attacker's application; the CRM provider later changed that default for all customers.

1.4. Qantas also detected, contained and responded to the incident promptly.

1.5. On balance, these factors supported the preliminary view that the evidence did not point to significant omissions or failures in Qantas's security measures.

However, the relevant counterfactual is whether Qantas could have identified and mitigated this attack chain through threat intelligence, targeted training, or other technical and organisational controls.

As the timeline in Appendix A shows, substantial public intelligence existed before the breach:

2.1. On 12 March 2025, Salesforce, publicly identified as the CRM, described the exact attack chain[1]: callers impersonated IT support and directed users to authorise malicious connected apps that exfiltrated Salesforce data. It also identified several compensating controls.

2.2. On 29 May, ACSC warned that social engineering targets staff able to grant system access and that vishing requests involving access changes should be independently verified.

2.3. On 4 June, GTIG reported UNC6040's repeated use of the same method, calling malicious connected-app authorisation a prevalent tactic[2].

The published report does not state whether Qantas monitored relevant threat intelligence, assessed its Salesforce exposure, or deployed targeted training or other mitigations before the incident.

The explanation below is an inference tightly constrained by the attacker, target, feature and tactic. Salesforce reported it 108 days before Qantas's breach; GTIG confirmed it 24 days prior:

3.1. The attackers targeted the widely used CRM *precisely because organisations like Qantas stored millions of valuable customer records within it*.

3.2. The attackers identified a CRM feature that let an end user authorise third-party access to the records, *making it an obvious target for a group that specialises in social engineering*.

3.3. The attackers chose vishing as an effective, reliable route to the vulnerability. *It looked legitimate, required no stolen credentials and evaded controls against credential theft*.

3.4. Therefore, the fact that the feature was legitimate or default-enabled does not make its abuse anomalous from a security perspective; *it is precisely the kind of feature these attackers would seek out*.

3.5. Therefore, using phone-based social engineering to induce this authorisation was a predictable attack path. *For these attackers, it was the most direct and efficient way to connect an unauthorised third-party application to the customer records*.

[1] Protect Your Salesforce Environment from Social Engineering Threats, Salesforce, 12 March 2025

[2] The Cost of a Call: From Voice Phishing to Data Extortion, Google Threat Intelligence Group (GTIG), 4 June 2025

The report relied partly on the view that standard training does not usually cover this specific tactic and that stronger RBAC would not have prevented it. These points informed its preliminary view that the evidence did not point to significant security failings and that Qantas could not reasonably have foreseen or prevented the breach. This reasoning may weaken the OAIC's position in future cases.

Every cyber-attack can be described at different levels of detail. At a broad level, this incident involved social engineering. At a narrower level, it involved vishing. At a narrower level again, it involved an impersonated IT support call that led an agent to authorise a third-party connection.

The more narrowly an attack is described, the more unusual it may appear.

If an uncommon subtype is treated as evidence against a significant omission or failure, future organisations may advance the same argument after many different breaches.

They could say their controls addressed the broader, more common risk, but that the attacker used an uncommon variation or abused an ordinary system function.

For high-profile cases, a Delphi process, or similar structured expert review, would give the OAIC a transparent, defensible basis for assessing normal practice and material departures from it. It would inform, not replace, the OAIC's statutory assessment.

The OAIC could draw on independent views from threat intelligence, offensive security, vendor-risk management, incident response, privacy, security awareness and risk management.

Their structured, anonymised judgments would provide an additional basis for deciding, where warranted, that an organisation reasonably managed a risk despite a breach.

A practical process would:

1. Give the panel a common, evidence-based case file
2. Collect anonymous initial judgments, reasons and confidence ratings
3. Share anonymised ranges, reasons and contrary evidence
4. Conduct a second round of assessment
5. Record the consensus, material disagreement, evidence relied on, and the OAIC's reasons for accepting or departing from it

If this approach is of interest, I recommend that the OAIC commission a research project to adapt it to regulatory decisions, test it against past cases, and publish the method and findings.

Applied to this case, such a process could have tested whether Qantas's threat-intelligence capability was adequate, whether the breach was reasonably foreseeable and capable of mitigation, and whether Qantas's practices were comparable with its peers.

Ultimately, this may not have changed your decision not to commence a CII, but it would have made that decision more defensible and reduced the risk of weakening the OAIC's position in future cases.

Kind regards,
Benjamin Mosse

Appendix A – Public threat intelligence and security guidance documented phone-based help-desk attacks for years and the exact Salesforce attack chain before the Qantas incident

The following cyber threat intelligence was publicly available before 28 June 2025. It does not establish that every named threat cluster was the same group.

It does show that phone-based social engineering against help desks, outsourced service providers and SaaS platforms was a well-documented risk. More importantly, Salesforce described the precise connected-app attack chain 108 days before the Qantas incident.

Date	Public Cyber Threat Intelligence	Relevance to the Qantas incident
1 December 2022	CrowdStrike reported a campaign, attributed with low confidence to Scattered Spider, targeting telecommunications and business-process outsourcing companies. Attackers impersonated IT staff by phone and directed employees to disclose credentials or run legitimate remote-access software.	Established phone-based impersonation and outsourced service providers as recognised attack vectors.
14 September 2023	Mandiant reported that UNC3944 persistently used phone-based social engineering, called help desks for password and MFA resets, targeted BPOs and used legitimate software during data-theft operations.	Established the combination of phone calls, support personnel, legitimate tools and large-scale data theft.
25 October 2023	Microsoft reported that Octo Tempest, which overlaps with Scattered Spider and UNC3944, researched target organisations and used phone calls to manipulate support personnel with administrative access.	Showed that attackers deliberately selected employees whose ordinary permissions could unlock valuable systems.
16 November 2023	The FBI and CISA warned that Scattered Spider targeted large companies and contracted IT help desks through voice communications and trusted relationships. They recommended testing and validating controls against the documented techniques.	Government guidance treated phone-based help-desk attacks as conduct organisations should test against.
8 February 2024	ASD's ACSC co-authored guidance on "living off the land". It explained that attackers abuse legitimate, trusted system features and that insecure default configurations can enable malicious activity. It recommended baselining behaviour, application allowlisting, detailed logging and controls over human-to-software interactions.	Directly supports your argument that a legitimate or default-enabled feature can be an attractive and foreseeable attack path requiring compensating controls.
9 April 2024	Unit 42 reported that Muddled Libra researched which applications and cloud providers an organisation used, social-engineered help desks and abused legitimate cloud functions for access and exfiltration.	Supported the attacker model advanced in this letter: identify useful software features and work backwards to an effective attack chain.
13 June 2024	Mandiant reported that UNC3944 used phone-based social engineering against corporate help desks and then accessed SaaS applications, including Salesforce. Its analysis included forensic recordings of call-centre attacks.	Connected help-desk social engineering directly with Salesforce and other SaaS environments.

Date	Public Cyber Threat Intelligence	Relevance to the Qantas incident
12 March 2025	Salesforce warned customers that attackers were impersonating IT support by phone and directing users to authorise malicious connected apps, sometimes modified versions of Data Loader, to exfiltrate Salesforce data. Salesforce recommended connected-app restrictions, application allowlisting, trusted IP ranges, limited Data Loader access and transaction monitoring.	Described the essential Qantas attack chain and possible compensating controls 108 days before the incident.
4 May 2025	The UK NCSC advised organisations to review help-desk authentication and password-reset processes and ensure they could rapidly consume and act on threat intelligence. It referred cautiously to suspected Scattered Spider activity.	Demonstrated that rapid operational use of threat intelligence was itself being presented as good practice.
6 May 2025	Mandiant published detailed hardening guidance for UNC3944. It identified large enterprises with extensive or outsourced help-desk functions as targets and recommended stronger identity verification and targeted awareness.	Applied directly to a major Australian organisation using an overseas contact centre.
16 May 2025	Unit 42 updated its Muddled Libra assessment to emphasise social engineering of end users and help desks, access through BPOs and training addressing suspicious phone and SMS contact.	Reinforced that non-email social engineering and outsourced support functions required specific attention.
29 May 2025	ASD's ACSC published guidance on social engineering. It warned that attackers use vishing to impersonate staff, target IT service desks and people able to grant system access, and persuade staff to alter system settings. It said sensitive requests and configuration changes should be independently verified.	Closely describes the human and technical elements of the Qantas attack, only 30 days before it occurred.
4 June 2025	Google Threat Intelligence reported that UNC6040 was repeatedly impersonating IT support by phone and persuading employees to authorise malicious connected apps in Salesforce for large-scale data theft.	Documented the same Salesforce attack method used against Qantas 24 days before the incident.
27 June 2025	The FBI warned that Scattered Spider was expanding its targeting to airlines and their third-party providers through employee and contractor impersonation.	Made the aviation sector and its service providers an express target, although only one day before the incident.

Attribution Note:

These actor names should not be treated as interchangeable. Microsoft describes Octo Tempest as overlapping with Scattered Spider and UNC3944. Google tracks UNC6040 separately and has said that similarities with actors associated with "The Com" may reflect related actors rather than a direct operational relationship. Google's June report also recorded attackers claiming an affiliation with ShinyHunters but did not establish that the names represented one organisation.

This timeline does not prove that Qantas received or reviewed these reports. It shows that the adequacy of Qantas's threat-intelligence capability and response is a material question not addressed in the published report.

Appendix B – A targeted warning Qantas could have sent

Salesforce described this attack method on 12 March 2025. ACSC warned about vishing-related access changes on 29 May, and GTIG reported the same Salesforce method on 4 June. Together, these reports could have provided the basis for a targeted warning to contact-centre agents, such as:

Team,

We are aware of fraudulent calls from people claiming to be IT support.

Do not follow a caller's instructions to open Salesforce, enter a connection code, or authorise a third-party application. End the call and verify the request through the approved IT support channel.

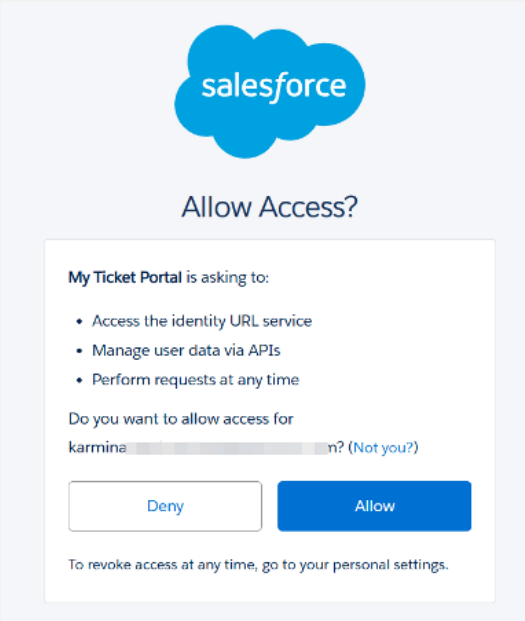

If you receive a suspicious call, end it and report it to the Security Team immediately.

Thank you,
Qantas Cyber Security

Image sourced from Cloud Protection for Salesforce.