


Breaching Boundaries: *Improving Data Breach Notifications in Australia*



THE UNIVERSITY OF
SYDNEY

Professor Jane Andrew, Associate Professor Max Baker and Dr Penelope Bowyer-Pont
Vol. 1 No. 1, *Data Breaches in Australian Organisations*. ISSN: 2982-1665.

The University of Sydney
Business School



We recognise and pay respect to the Elders and communities – past, present, and emerging – of the lands that the University of Sydney’s campuses stand on. For thousands of years they have shared and exchanged knowledges across innumerable generations for the benefit of all.

Authors:

Professor Jane Andrew, Associate Professor Max Baker and Dr Penelope Bowyer-Pont,
The University of Sydney Business School

Acknowledgments:

This project was funded by an Australian Research Council Discovery Grant.

July 2024

www.doi.org/10.25910/psq3-q365

Vol. 1 No. 1, *Data Breaches in Australian Organisations*. ISSN: 2982-1665.

Contents

Executive summary	2
Key findings and recommendations	3
Background: The Australian context	8
1. Definitional challenges: 'Likely to cause serious harm'?	10
2. Incentive design challenges: Ransoms, fines, and workplace culture	12
3. Regulation: More resourcing and powers?	14
4. Notifications: too much or not enough?	16
5. Third-party vendors: who notifies?	18
6. Data retention: how long to store data?	22
7. Production data testing: unnecessarily increasing risks?	24
8. Not-for-profits: unique challenges	26
9. Cyber experts: more than a niche skill?	28
10. Best practice: optimisation through industry collaboration?	29
Conclusion	31
References	32
List of acronyms	34

Executive summary

Data breaches occur when personal information is accessed, disclosed without authorisation, or lost (OAIC, 2024c). The World Economic Forum (2024) reports that instances of compromised data increased globally by 72% between 2022 and 2023. In Australia, data breaches are rising in both size and frequency, with a significant impact on individuals and organisations. For individuals, the consequences include compromised privacy, ‘serious harm’ (including, but not limited to, psychological, reputational, and financial harm) and identity theft. For organisations, breaches are costly and sometimes disastrous.

This report presents the preliminary findings from one component of a larger project, funded by the Australian Research Council, examining organisational data breach disclosure practices in Australia. This part of the project sought to understand how organisations navigate the increasingly important and challenging areas of information security, privacy, data breaches, and breach disclosure and notification. We conducted in-depth, semi-structured interviews with 50 senior personnel from organisations ranging from for-profit entities (including Australia’s big four banks and major insurance and superannuation providers) to some of Australia’s largest not-for-profits, as well as key government departments and agencies at both a state and federal level. We interviewed senior executives and personnel working in information security, privacy, cyber security, data management, risk, and compliance roles.

We wanted to understand the current state of the field, overall, and how organisations and their staff manage the increasingly complex challenge of protecting the personal information of customers, clients, and consumers in the digital age. Importantly, the study explores how cyber ecosystems impact data breach notifications within organisations, with a view to improving practice.

In this report we present ten key findings from our interview analysis. For each finding we suggest recommendation(s) that, if implemented, support an improved disclosure regime and can help guide organisational best practice as data breach notification assessments are made. This can then help to inform any subsequent notifications to the regulator and impacted individuals.

Key findings and recommendations

Finding



1. Definitional challenges: ‘likely to cause serious harm’?

In practice, definitions of ‘data breach’ and ‘serious harm’ are highly subjective and open to interpretation. Organisational leaders make decisions on a case-by-case basis about what constitutes a breach and whether to notify. Because of this, we found a spectrum of reporting. We also found that not-for-profits (NFPs) are more likely than large commercial entities to have a lower threshold for notification.

Recommendation



Organisational leaders would benefit from more granular and case-based guidance from the Office of the Australian Information Commissioner (OAIC) on interpreting different forms of harm. The United Kingdom’s Information Commissioner’s Office (ICO) was cited by many interviewees as providing better guidance, including a self-assessment tool to determine the need to notify and detailed case studies to help guide assessments. Given this support can be provided at low cost, the OAIC should develop a similar self-assessment tool with built-in case examples to guide organisations at each stage of the notifiable data breach (NDB) assessment process.

2. Incentive design challenges: ransoms, fines and workplace culture

Ambiguity around what constitutes a data breach that is likely to result in serious harm to an individual can allow organisations to downplay the severity of breaches. Furthermore, businesses often choose to pay ransoms, thereby mitigating notification requirements, while also avoiding scrutiny from the regulator, the media, and stakeholders. Paying ransoms further incentivises threat actors.

Organisational culture has an impact on the thresholds for data breach notifications. When data breaches occur, employees may be reluctant to inform their managers for fear of reprimand. At the corporate executive level, a pervasive fear of reputational harm means that organisations may downplay the seriousness of harm associated with breaches.

There needs to be meaningful financial consequences associated with data breaches. Organisations should face impactful fines for either failing to protect personal information properly or for not responding to breaches adequately. This includes failures to provide appropriate notifications to the regulator and individuals. These fines need to be large enough to incentivise organisations to invest properly in cyber security and to report breaches when they occur.

Organisational cultures need to support robust internal data breach reporting practices. The reporting of incidents must be aligned to the interests of the organisation. Policies that outline employee responsibilities and internal processes need to be developed in organisations of all sizes to counter fears of blame and shame.



3. Regulation: more resourcing and powers?

An adequately funded and equipped regulator can respond promptly, enhancing the effectiveness of breach notification procedures. Interviewees believe the OAIC lacks the necessary authority and resources proportional to the complexity of preserving privacy and safeguarding personal data within an evolving digital landscape.

In the event of a breach, organisations are obligated to report to various entities beyond the OAIC. This process is cumbersome, perplexing, and likely to increase mistakes and miscommunication as it involves engaging with multiple parties.



The regulator's authority must be enhanced to address effectively the increasing challenge of safeguarding personal information and minimising the repercussions of data breaches. Adequate resources should be allocated to the OAIC to ensure it has forward-looking capabilities. To support an effective and timely notifications regime, the OAIC needs additional public facing guidance and a well-resourced response team.

Additionally, organisational leaders would benefit from simplified and streamlined notification responsibilities. Having the OAIC as the sole recipient of data breach notifications would be simpler and more straightforward. It could then be responsible for notifying other relevant agencies and bodies, including the Australian Tax Office (ATO), Home Affairs, the Australian Cyber Security Centre (ACSC), the Australian Prudential Regulatory Authority (APRA), and law enforcement.

4. Notifications: too much or not enough?

Participants recognised that data breach notifications themselves may cause harm, with serious psychological and emotional harm potentially caused to an impacted individual learning their personal information has been compromised.

Given the significant implications associated with both notifying and not notifying, organisations need to balance complex and often competing considerations, noting that standard notification letters can cause distress and confusion.

Impacted individuals need to be provided with contextual details explaining why their data was being held, the nature of the breach, and the next steps that need to be taken – by both the organisation and the individual – to minimise potential harm. Notifications should include more information about the specific risks associated with each type of data involved in breaches (e.g., Medicare number, drivers' licence, home address, etc.)

In addition, organisations need to tailor their notification processes for vulnerable individuals, such as people with disabilities or sensitive health issues, and victims of domestic, family, and sexual violence. For these groups, notification is rarely straightforward.

5. Third-party vendors: who notifies?

Reliance on third-party vendors or suppliers has become commonplace for organisations across various industries. While often essential, outsourcing heightens the risk of data breaches. All interviewees expressed concerns about the adequacy of protection of sensitive information by their vendors. The OAIC (2024d) issued a statement in February 2024 highlighting the ongoing surge in multi-party breaches, primarily attributed to cloud or software providers. Multi-party breaches create challenges in determining the organisation responsible for data breach notifications in complex, often cross-national supply chains.

Enhanced oversight of third-party vendors that extends beyond conventional vendor management checklists is imperative. Australia should adopt privacy legislation akin to Europe's General Data Protection Regulation (GDPR), which specifically and legally delineates the data protection and data breach notification responsibilities of 'data controllers' (organisations that collect data) and 'data processors' (those that manage or host data on behalf of another organisation). With the Privacy Act 1988 (Cth) undergoing review, interviewees support a transition toward this framework. However, swift implementation is essential.

Currently, the OAIC (2024e) suggests that the responsibility for notifying impacted individuals of breaches rests with the organisation with the most direct relationship to the individuals at risk of harm. This needs to be legislated in the amended Privacy Act to create clarity and consistency.



6. Data retention: how long to store data?

Unnecessary data retention greatly increases the likelihood of breaches. Breaches are frequently caused by vulnerabilities in legacy systems that organisations have not decommissioned. It is very common for organisations, including large financial institutions, to be unsure of exactly what data they hold and where it is, making it impossible to secure. There is considerable confusion among senior executives about their data retention and destruction obligations, creating a situation where data is retained ‘just in case’.



More comprehensive compliance assessments and data audits would help to ensure organisations do not retain data simply because disposal or decommissioning is difficult, time consuming, and costly. Australians would benefit from the implementation of the ‘right to erasure’ (akin to the GDPR), which gives ‘data subjects’ (individuals) the right to ‘be forgotten’ by organisations and have their personal information erased.

For the Notifiable Data Breaches Scheme (NDBS) to be managed effectively for both organisations and individuals, retention rules must be clarified and streamlined.

The OAIC could mandate that breaches occurring due to vulnerabilities in legacy systems require organisations to explain (as part of their notification) to impacted individuals why their personal information was held past the mandatory data retention period. This would encourage organisations to dedicate more time and resources to properly decommissioning legacy systems, which would better protect all Australians.

7. Production data testing: unnecessarily increasing risks?

Senior executives are worried about the common practice of production data testing. The use of real, personally identifiable information for the development and testing of products creates a significant risk of breaches because testing often takes place in a less secure environment. Many interviewees expressed concern about this practice, which appeared to be particularly common in the financial services sector because of the competitive speed-to-market culture.

Australia’s privacy regulator must ensure that organisations are adhering to Australian Privacy Principle 11 (APP 11) of the Privacy Act. Testing with personally identifiable information that has not been deidentified risks individuals’ data being breached and is not consistent with APP 11. As with other recommendations, addressing this requires a stronger regulator that can enforce the Privacy Act.

If a breach is due to production data testing, the OAIC should mandate that this information be included in the notification to both the regulator and impacted individuals. This would incentivise organisations to properly deidentify data before using it for testing purposes.

8. Not-for-profits: unique challenges

NFPs working with vulnerable people hold highly sensitive personal information, including health data. NFPs often lack adequate resources to safeguard health data effectively in comparison to commercial entities. This means that there are fewer resources for helping vulnerable NFP clients when their data is compromised. This issue has become particularly concerning as threat actors increasingly target environments with (i) health data and (ii) low security and high levels of personal data.

NFPs would benefit from dedicated funding and resources to enable them to better protect highly sensitive health data. More broadly, increased focus on the security of data held by NFPs is needed, given that the current focus is skewed towards the security of financial data in large commercial entities. This neglects the ‘serious psychological ... or reputational harm’ (OAIC, 2024a) that can result from breaches of highly sensitive health data.

NFPs have a lower threshold for notifications, and notifications often need to be carefully tailored for people with complex needs, such as those with a disability or the elderly. For this reason, NFPs also need additional resources to ensure notifications are fit for purpose.



9. Cyber experts: more than a niche skill?

Australia faces a significant skills shortfall in information security and data protection. There are currently not enough cyber security professionals with the right skills and experience to meet increasing demand, a problem made worse by the speed at which threat actor capabilities are developing. This has implications for an organisation's capacity to assess the level of harm associated with a breach and the appropriate notification pathway. Indeed, attracting cyber security professionals is particularly difficult for less-resourced small and medium-sized businesses and NFPs.



Government investment is needed to upskill existing professionals and train many more. Subsidised schemes encouraging more people to train in the crucial areas of cyber security and data protection can address the shortfall over the coming decade.

Universities, too, need to expand and promote their cyber security course offerings and consider hiring more academics in this space.

Organisations could focus more on upskilling their existing employees and expanding internship programs to help new graduates gain valuable workplace experience. This approach would ensure that organisations have the expertise needed to conduct timely and evidence-based assessments of harm, enabling them to notify affected individuals appropriately. This, in turn, supports the implementation of effective risk mitigation strategies for those impacted.

10. Best practice: optimisation through industry collaboration?

Data breaches pose a collective threat, and safeguarding data requires collective effort. Developing best practices for communicating breaches requires collaboration. Every individual and organisation is susceptible to risk, underscoring the importance of mutual cooperation. Our interviewees stressed the significance and benefits of inter-organisational sharing of insights gained from breaches and post-breach notifications strategies, particularly communications with impacted individuals.

All organisations would benefit from more structured opportunities to share experiences and learnings from breaches and subsequent notifications. The Australian Signals Directorate (ASD) runs the Australian Cyber Security Partnership Program (ASCPP), which helps to facilitate the sharing of knowledge and expertise to collectively reinforce cyber security resilience across Australia. The ASCPP's work must be well resourced and ongoing to ensure best practices are shared across industries and sectors, and that organisations' notification practices are informed by current best practice.

"My point has always been it is subject to *interpretation*, and this is where it is very unclear what is serious harm. What [one person] might consider serious harm versus what I might consider serious harm could be quite *different*."

(Participant 7)

Background: The Australian context

In the second half of 2022 and early 2023, millions of Australians were impacted by several high-profile data breaches, including that of Optus, Medibank, and Latitude Financial. These followed other major breaches, such as that experienced by the Australian National University in November 2018 (impacting 200,000 students), Service NSW in 2020 (104,000 people), and ProctorU in 2022 (440,000 people).

These events have focused attention on the omnipresent threat to personal information from cyber security incidents. In August 2023, Australia's information and privacy regulator, the OAIC, released the findings of the Australian Community Attitudes to Privacy Survey, which found that 75% of Australians consider data breaches to be one of the primary risks to their privacy (OAIC, 2023a). Attempts to prevent data breaches and protect privacy are outpaced by the increasing sophistication of threat actor strategies. In the second half of 2023 prominent cyber security company SurfShark released a report from its Data Breach World Map (SurfShark, 2023) detailing the surge in data breaches globally and showing Australia as ranking fifth in the world for density of breaches (number of compromised accounts per 1,000 residents). The World Economic Forum (2022) now recognises cyber risk as *'the most immediate and financially material sustainability risk that organisations face today'*.

The escalating threat of (and focus on) data breaches takes place alongside a major overhaul of Australia's *Privacy Act 1988* (Cth), which began in 2020. In February 2023 the Federal Government said that it would accept 106 reform proposals to bring the Privacy Act into the digital age. The coming changes aim to upgrade organisational data protection, increasing transparency and improving individuals' control over their information. Regulation is attempting to catch up to cyber threats. In 2022 the Australian Securities and Investment Commission (ASIC) took legal action against an Australian financial services licensee, RI Advice, for failing to adequately manage its cyber security risk. The Federal Court found that the financial services company breached its obligations under the *Corporations Act 2001* (Cth) by failing to invest in the protection of customer data. This case represents the first time an Australian financial institution has been found to be in breach of the Corporations Act due to cyber security-related conduct. It is also significant because it signals ASIC's willingness to hold directors legally accountable for inadequate cyber risk management. Despite these key developments, serious data breaches caused by both human error and by malicious actors continue to occur regularly in Australia.

In February 2018, the OAIC introduced the NDBS, which requires entities bound by the *Privacy Act* to notify both the regulator and impacted individuals of breaches ‘likely to cause serious harm’. The most recent report of the NDBS was released in February 2024 and covers the July–December 2023 period, during which the OAIC received 483 notifications of breaches, an increase of 19% (from 407) compared with January–June 2023 period. Of these breaches, 67% were the result of malicious criminal actors and the health and finance sectors were the top reporters of breaches, followed by the insurance and retail industries, then the Australian Government. Contact information is the type of personal information most compromised in breaches, followed by identity and health information (OAIC, 2024d). In this most recent reporting period, health-related data overtook financial information as the third most common type of personal information affected by breaches. For malicious criminal actors, health information is far more valuable on the dark web than credit card details, offering a possible explanation for the significant rise in breaches of this kind. Credit card details can be easily changed, while health data usually encompasses all of an individual’s personally identifiable information and cannot be altered.

Australia’s privacy regulator requires entities to take reasonable steps to complete an assessment and notify of a breach within 30 days of an incident. However, Fell et al. (2023) report that many significant breaches are not disclosed to the regulator because organisations assess the potential harm caused by such breaches internally. This means that organisational leaders largely set their own thresholds around what constitutes ‘serious harm’. The subjective nature of understanding and predicting ‘harm’ resulting from breaches remains a challenge and will be discussed in the next section of this report.

Many *significant* breaches are not disclosed to the regulator because organisations assess the potential *harm* caused by such breaches internally. This means that organisational leaders largely set their own *thresholds* around what constitutes ‘serious harm’.



Definitional challenges: 'Likely to cause serious harm'?

One of the most significant challenges our interviewees reported was how to accurately define and effectively respond to a notifiable data breach. The Privacy Act defines a data breach as: *'an unauthorised access or disclosure of personal information, or loss of personal information'*. Under the OAIC's NDBS a breach must be reported to the regulator if it is *'... likely to result in serious harm to any of the individuals to whom the information relates'* (OAIC, 2024a). While the definition appears relatively straightforward, the terms 'unauthorised access', 'personal information', 'disclosure', 'loss', 'likely', and 'serious harm' all require interpretation. In larger organisations, including large financial institutions, internal general counsel is usually responsible for providing legal advice about whether incidents constitute notifiable data breaches. Small to medium-sized enterprises (SMEs) often seek external legal advice, but some organisations do not have privacy officers or the resources to engage external consultants. In these cases, it falls to senior management to interpret the OAIC guidance and make judgements on a case-by-case basis as to whether a breach is *'likely to result in serious harm'*.

Cyber incidents and events happen all the time and 'unauthorised access' is very common. According to the definition in the Privacy Act, if an employee looks at another employee's computer screen and views personal information they would not otherwise see, a breach has, technically, occurred. Similarly, if someone accidentally sends an email containing personal information to the wrong person within an organisation, a breach has occurred. In practice, however, the labelling of incidents as data breaches is often not straightforward. We discovered significant variations in how senior personnel across different organisations interpret breach definitions and, consequently, the necessary actions that follow, including reporting to regulatory bodies.

1.1 Data breaches

We found that, in practice, what constitutes a data breach is understood along a spectrum from any unauthorised access at one end and data exfiltration at the other. The former refers to gaining access without proper authorisation, while the latter occurs when data actually leaves the organisation, often as a result of deliberate theft. We identified a noticeable difference between the for-profit and NFP sectors, with a senior data professional working at a major Australian NFP telling us a breach is:

... anything where information is accessed by the wrong person ... or given to the wrong person ... It's really hard to decide for another individual what's harmed them and what's not, so we lean towards transparency.
(Participant 5).

Overall, our interviewees from NFPs tended to adopt a more cautious and rigid definition of what constitutes a 'data breach' and many said that they would prefer to 'over-report'. In contrast, a data professional working at a large ASX-listed company responsible for critical infrastructure told us:

It really needs to be that there is confirmed loss of data. Just because someone had access to a system doesn't necessarily mean that the data left the system ... [but] that's hard to prove sometimes one way or the other, depending on the level of logging and monitoring within that environment.
(Participant 1)

Thus, if we view organisations as sitting on a continuum with 'only reporting when necessary' at one end and 'reporting just in case' at the other end, then commercial entities tend to lie more towards the former, while NFPs are more towards the latter.

1.2 Harm

Harm is central to the definition of notifiable data breaches. On its own it is understood in a variety of ways by experts. Sometimes it is obvious and straightforward that a breach is likely to result in serious harm. More often, however, this assessment is difficult. A senior cyber security professional working for an ASX top-500 company told us:

My point has always been it is subject to interpretation, and this is where it is very unclear what is serious harm. What [one person] might consider serious harm versus what I might consider serious harm could be quite different. (Participant 7)

When assessing potential harm following a breach, timing is critical. Assessment is dependent on available information, which often emerges gradually over the course of several days, and sometimes much longer, while a forensic investigation is being undertaken. Organisations must strike a balance between timely notification to affected individuals and the need to gather crucial information about the incident, which takes time. An information security specialist, who works for a major NFP and aged care provider, told us:

We are making those harm assessments all the time. For the vast majority of them, there's no harm or very little harm. The problem, though, is that you are working on a limited data set so you might make this assessment based on the information that you have at the time that there was no harm, but then it turns out, in the fullness of time that actually that resulted in somebody being on your internal network. (Participant 8).

Overall, our research found that senior personnel face challenges when interpreting the OAIC's guidance on what constitutes 'likely' and 'serious harm'. Organisations must balance timely notification to affected individuals with gathering essential information about the incident, a process requiring time and precision. Participants expressed concerns about the complexity of determining harm, given potentially unknown additional factors in the life of an individual. This can mean that a seemingly benign breach, for example, that includes peoples' home addresses, could result in serious harm. This is particularly the case for domestic and family violence, which was referred to by multiple interviewees, as there have been cases where a domestic violence victim's safety has been compromised when a data breach has revealed their address to a former spouse. Our research participants frequently described the challenges associated with these nuanced scenarios.



Recommendation One:

Organisational leaders would benefit from more granular and case-based guidance from the OAIC on interpreting different forms of harm. The United Kingdom's ICO was cited by many interviewees as providing better guidance, including a self-assessment tool and detailed case studies to help guide assessments. Given this support can be provided at low cost, participants said they would benefit from a similar self-assessment tool with built-in case examples to guide organisations at each stage of the NDB assessment process.

Incentive design challenges: Ransoms, fines, and workplace culture

Our research found that organisations are not incentivised to report breaches. The ambiguity around the definition of key terms provided by the OAIC (see section 1) allows organisations to downplay the severity of breaches if it is in their interest not to report. In practice, this means that organisations need to determine what should be considered a ‘data breach’ and what is considered ‘serious harm’ and, therefore, decide whether to report. One information security specialist told us:

When I see the definition of it, I see that it's not very clear ... I find a lot of greyness in that definition. It's not really black or white. It doesn't really clearly articulate and say what needs to be disclosed ... it's very subjective ... That's where I feel there is a clear gap in terms of what needs to be reported and it is left, to some extent, to organisations, really, to decide, to interpret and apply their interpretation and decide what to report and what not to ... it's quite fuzzy right now as to what is reportable and what is not. (Participant 7)

Many interviewees had previously worked in various senior corporate roles and were able to comment on the data breach disclosure practices of their past workplaces. Under-reporting, we were told, is partly driven by an organisational culture that seeks to avoid embarrassment wherever possible. As one interviewee told us: *‘It's one of the biggest problems. The culture of [avoiding] embarrassment at the corporate level is huge’* (Participant 11).

We interviewed a former senior public servant, who previously worked in a lead role for the OAIC. This interviewee emphasised the significant problem of under-reporting, pointing out that probably only 10% of data breaches are reported:

... the proportion of data breaches in Australia ... that would pass a test of ‘likely’ and ‘serious’, [that are actually] notified is, in my view, likely to be 10% or less. (Participant 10)

Partly, this is due to the level of subjective interpretation required to enact the law at the organisational level, and partly a lack of incentives to report. However, interviewees also told us that there is not enough awareness about the Privacy Act and the NDBS. The same former senior public servant told us:

I think you could actually do a survey of all the companies in Australia to find out what proportion of them are even aware of the existence of the Privacy Act. Of that proportion, what subset is aware of the notifiable data breaches requirement and what subset of that actually knows what the requirements are. (Participant 10)

In addition, a leading cause of day-to-day data breaches is workplace cultures that are unforgiving of human error. Participants told us that breaches are sometimes not reported because employees are reluctant to tell their managers about data incidents or breaches for fear of reprimand. One interviewee described it in this way:

... it has a lot to do with organisational culture ... often people get a little bit nervous and scared ... if somebody inadvertently sent a sensitive email outside the organisation to a recipient who was not supposed to receive it ... it is the organisational culture that drives a lot of it ... how individuals are treated when something like this happens, whether they are reprimanded or assisted and supported ... (Participant 7)

2.1 Ransoms

Ransom payments are another significant problem. Participants told us that companies are willing to pay ransoms to threat actors to resolve issues. These incidents then usually go unreported because organisations believe that by paying the ransom demands of threat actors, they have resolved the issue and the regulator or impacted individuals do not need to be informed. However, this is a very risky and dangerous practice, allowing those companies to avoid scrutiny from the regulator, the media, and stakeholders.

The ACSC warns entities against paying ransom demands. Doing so not only allows the incident to go unreported and uninvestigated, it also incentivises further attacks from threat actors. Ransomware attacks currently cost the Australian economy AUD2.59 billion annually (*Australian Cyber Security*, 2024). A survey by Sydney-based accounting firm McGrath Nicol of Australian businesses in 2023 found that 73% of businesses that experienced a ransomware attack between 2018 and 2023 chose to pay the ransom. The average ransom paid was AUD1.03 million, with business leaders saying they would willingly pay an average of AUD1.32 million to resolve a ransomware attack (McGrath Nicol, 2023, p. 1).

Our research findings confirm that the paying of ransoms is a common practice in Australian organisations. One of our interviewees was an information security specialist at a global cyber security consultancy, who had worked in various corporate roles over the course of their career. At one point they worked for an organisation that experienced a massive data breach, caused by a malicious attacker, with devastating consequences. Referring to the paying of ransoms in their former workplaces, this interviewee told us:

They just pay the money, get the keys and pretend it never happened. They're happening every day of the week ... They have a meeting in a coffee shop down the road and go, 'What the hell just happened? Alright, let's not tell anyone about this and let's just move on. Pay the ransom. Let's move on.' (Participant 9)

Ransom payments take place despite the Federal Government urging businesses not to pay. If ransoms are paid, breaches are unreported, which means that the regulator and the individuals whose data has been compromised are not notified. This is a huge problem because the regulator and impacted individuals cannot respond to breaches if they are not made aware of them. In addition, when threat actors receive a ransom payment, they may still publish exfiltrated data on the dark web, despite promising the organisation otherwise, amounting to double extortion. At the unveiling of the Australian Cyber Security Strategy 2023–2030, Home Affairs and Cyber Security Minister, Clare O'Neil, said:

'Every time a ransom is paid, we are feeding the cyber crime problem' (Australian Financial Review, 2023). The Government has indicated its intention to ban the paying of ransoms, making this an illegal act, in the next two years.

The overall finding from our research is that there is considerable under-reporting, which occurs, in part, because organisations determine for themselves whether a breach meets the threshold required to report. In addition, our research points to two other significant drivers of under-reporting. First, the paying of ransoms and second, organisational cultures that make it difficult for employees to speak up about breaches for fear of consequences. A study in 2019 by the Information Systems Audit and Control Association (ISACA), a global organisation for information technology and cyber security professionals, found that under-reporting, even when disclosure is legally mandated, is the norm (ISACA, 2019). Our findings confirm this ongoing trend.



Recommendation Two:

There must be meaningful financial consequences associated with data breaches. Organisations should face impactful fines for either failing to protect personal information properly or for not responding to breaches adequately. This includes failures to provide appropriate notifications to the regulator and individuals. These fines need to be large enough to incentivise organisations to invest properly in cyber security and to report breaches in a timely fashion when they occur.

Organisational cultures need to support robust internal data breach reporting practices. To achieve this, the reporting of incidents must be aligned to the interests of the organisation. Policies that outline employee responsibilities and internal processes need to be developed in organisations of all sizes to counter fears of blaming and shaming.

Regulation: More resourcing and powers?

Our interviewees believed the OAIC needs increased powers and resources in the fight to prevent data breaches and protect personal information, and that this would help improve the overall quality of notifications. Interviewees described several inter-related problems. First, the need for more and detailed instruction about what constitutes a breach and what constitutes harm, and the actions that need to be taken (see our section 2 recommendation). Research participants repeatedly told us that clearer guidance and more information would overcome some of the ambiguity surrounding notification and reporting. One interviewee, who works for one of Australia's large financial institutions, told us:

The OAIC's got guidance... but it's all written in that regulatory way where they don't want any definitiveness in case they get criticised, so still, even with guidance, you're trying to work your way through, what are they actually saying? (Participant 4)

Second, interviewees said that the OAIC needs to be better resourced and should have greater powers to impose and enforce fines. Organisations need to face impactful financial consequences for not investing properly in information security, or for not responding appropriately to breach incidents, which includes the provision of appropriate notifications. Interviewees told us that impactful fines are important to drive appropriate notification. The same participant told us:

The OAIC is renowned as a regulator that is just not doing its job... That's what we have here, a very weak regulator... You need a strong regulator who's gonna come in, sweep through your business give you a whopping great fine... that's when you'll start to see change... a stronger regulator who is a lot more definitive in their position makes you sit up a bit straighter – there's more impetus isn't there? There's more focus on doing the right thing. (Participant 4)

An information security specialist working for an ASX-listed entity expressed the same sentiment:

... they need to pull their socks up and do the enforcement. What's the point of having a policy if it's not enforced? ... I think the regulators just need to follow through ... come in, knock on the door and say 'we want to check your [security] environment' ... They really need to catch people. What's the point of having speeding signs and cameras and that, if you don't give anyone a ticket? (Participant 1)

Participants indicated concerns about what they viewed as insufficient resourcing of the regulator. They told us that it can be difficult to get a proper response when contacting the regulator. Some interviewees said that when they tried to contact the OAIC to either notify them of a breach or to ask a question, the regulator's response was inadequate. Many of the organisational leaders we spoke to said that their questions were often handled by junior OAIC staffers, who are often ill equipped to respond to such enquiries. We were also told that the OAIC does not always respond or follow-up when organisations make voluntary notifications about breaches and, as a result, the opportunity to reinforce good organisational behaviour was missed. Most suggested this was to do with resourcing, with one of our interviewees from an international NFP telling us:

At the moment, you make a privacy complaint to the regulator... within a year, they might pass that on to the organisation... that's not ideal, but that's the kind of funding that they've currently had... the people are good people, but they just don't have the money to do the large scale investigations, to do audits and things. (Participant 15)

Interviewees also expressed concern about what they viewed as the proliferation of distinct regulatory bodies, in addition to the OAIC, that demand data breach notifications. In certain situations, industry-specific bodies must be notified. For instance, in healthcare, the Australian Health Practitioner Regulation Agency (AHPRA) serves as a relevant authority. Many organisations are also required to report to the ATO, the ACSC, AUSTRAC (Australian Transaction Reports and Analysis Centre, monitoring financial crime), ASIC, law-enforcement bodies, the entity's own insurance provider, and the state or territory-level information commissioner. Research participants said that notifying many separate bodies is time consuming, and complicated, and that the potential for error is greater when there are multiple points of contact with different bodies. We interviewed a privacy expert for an ASX top 30 entity, who had prior extensive experience working for a multinational corporation, giving them a unique perspective that enabled insightful comparisons between the regulatory frameworks in Australia and Europe. They told us:

It would be better if you didn't have to go out to three or four different bodies to notify about the same breach ... it would be useful to be able to notify in one place and then have the Privacy Commissioner be made aware, law enforcement, the cyber center, because that could all be done potentially at the same time ... rather than having to put out those notifications separately. (Participant 6)

Other interviewees echoed a similar sentiment. The proliferation of regulatory bodies involved in the data breach notification landscape has led to confusion and overwhelm. When an organisation faces a breach, it must allocate resources and time to fulfill multiple notification requirements, diverting attention from thorough forensic investigation and harm prevention. Recognising the need for streamlined notification and cyber security regulation, the Australian Government established the National Office for Cyber Security. In February 2024, Lieutenant General Michelle McGuinness assumed the role of National Cyber Security Coordinator.

Recommendation Three:

There is a significant regulatory problem with data breach disclosure in Australia. At present, the OAIC lacks the necessary power and resources to adequately safeguard privacy and protect information in our ever evolving and vulnerable digital landscape. Australia would benefit from a stronger regulator that can meet the increasing challenge of protecting personal information and mitigating harm from breaches through timely and meaningful notifications. The OAIC needs significantly increased power, funding and capacity.

In addition, the notification process needs to be streamlined. Having one primary body responsible for receiving reports of a breach, which then communicates this to other relevant agencies or bodies, would be simpler and less time consuming for organisations. This approach also leaves less room for error and miscommunication. The OAIC currently co-chairs the Cyber Security Regulator Network (CSRN), *'a forum for Australian regulators to work together to understand, respond to and share information about cyber security risks and incidents. The CSRN works to reduce duplication or gaps in regulatory responses, so that regulatory activities are effective and efficient'* (OAIC, 2024d). The CSRN, or multiple reporting bodies, may not be needed if the Government upscaled and legally empowered one central regulatory body.



Notifications: Too much or not enough?

Our research identified significant problems with the ways individuals are notified when their personal information is compromised following a data breach. First, standardised notification letters put the burden of remediation on the individual. This is particularly problematic when the compromised data comes from a legacy system, that is, the organisation has been keeping data for longer than necessary (see data retention discussion in section 7). Second, standard notifications generally provide no explanation as to why the data was being held by the entity. Lastly, standard notifications presuppose that everyone possesses the ability to respond and take the necessary 'next steps' to safeguard themselves.

Organisations need to consider harm from breaches, as well as the potential harm associated with notification. In some cases, there is the potential for harm resulting from an impacted individual being told that their personal information has been compromised. In these cases, the psychological and emotional harm of being notified can be greater than any potential financial or reputational harm. Many interviewees described this dilemma.

To mitigate the potential harm caused by notifying, one state government department has developed a process to protect vulnerable people. They added high risk markers in their database for individuals on their system who are likely to need extra support if they are notified that their personal information has been breached. The department sought the advice of clinical psychologists, who recommended that a response pathway be built into the system that would allow for impacted individuals to receive more information about the context of the breach and why their personal information was being held, as well as some additional information about actions they might need to take.

"...the vast majority of our breaches are involving people *suffering* from family or domestic violence. That's where we can't fix the *problem*."

(Participant 4)



A senior government department manager told us ‘We found that providing the contextual information about what we were doing with that information [why the information was held] went a long way to try and alleviate the concerns’ (Participant 16). Impacted individuals can better comprehend the situation and take necessary precautions when they receive detailed information about a breach involving their personal data.

Many interviewees expressed concerns about notification practices and their potential impacts on individuals in the context of domestic violence. In some instances, an organisation has inadvertently sent a breach notification letter to an individual’s outdated address making it possible for an ex-spouse to learn of the individual’s current residence. We learned that in most organisations, the notification processes would not be sufficiently attuned to the vulnerabilities of breached individuals. One interviewee said:

... typically, the vast majority of our breaches are involving people suffering from family or domestic violence. That’s where we can’t fix the problem. Somehow, we’ve sent a letter that should have gone to Mrs Smith, and it went to Mr Smith, and that person was able to then track the other person. Once that person suffers some sort of psychological harm or mental distress, we can’t – you can’t fix that type of stuff. You can try and relocate people. You can pay for security on their homes. You can take a number of different steps that we have at our disposal, but by and large, you can’t fix that type of harm.
(Participant 4)

There have also been cases where an organisation had to notify someone that their HIV positive status had been exposed in a data breach. The notification letter to inform the impacted individual included reference to their HIV positive status. That notification letter was then sent to the individual’s address but was opened by a family member or someone other than the intended recipient. In this case, the person’s highly sensitive health information was compromised by both the breach and by the notification of the breach. Describing a case like this, one participant reflected:

How do you tangibly remediate the fact that your HIV status has now been published? Something you’ve held dear as a secret, and rightly so, and is protected by legislation. That’s now out in the wild.
(Participant 16)

The ways that organisations mitigate and remediate are also important. It is often the choice of actions taken following a breach that determines the likelihood of resulting harm. If an organisation can take action to address a breach after it occurs, they often choose not to notify individuals or report to the regulator, because the potential for harm has been mitigated. This is allowed for in the Privacy Act, which requires that a breach be reported if an entity has been ‘unable to prevent the likely risk of serious harm with remedial action’. In this way, mitigation and remediation are factors that help organisations in their decision making around reporting and disclosure following breach incidents, and they can minimise the risks to individuals associated with notifications.



Recommendation Four:

Australians would benefit from improved data breach notification practices. Standardised notification letters can cause distress and confusion. Most notifications also burden the individual with remediation work and this burden needs to be lifted. An individual impacted by a breach would benefit from the provision of contextual information that explains the nature of the breach, why their data was being held, and the next steps that need to be taken – by both the organisation and the individual – to minimise potential harm. Notifications should include more information about the specific risks associated with each of the types of data involved in breaches, for example, the risks associated with the breach of a Medicare number, a drivers’ licence, or a home address. Organisations need support and guidance to tailor their notification processes for vulnerable individuals and should develop best practice frameworks for appropriate notification for people with, for example, disabilities or sensitive health conditions and victims of domestic, family, and sexual violence.

Third-party vendors: Who notifies?

Organisations pay other organisations for specific services and to carry out day-to-day operational and managerial tasks. In doing so, primary organisations (or ‘data controllers’, as they are referred to in Europe’s GDPR legislation) pass on the personal information of customers and clients to these third-party vendors (or ‘data processors’). The outsourcing of tasks and the transfer of data to third-party vendors is an increasingly common and often necessary organisational practice. In the digital world, however, the use of third parties creates many additional security vulnerabilities and risks. An organisation can have the strongest possible internal data protection practices and an excellent cyber security framework, but if the vendors in their supply chain do not have the same standard of security and information protection, data can be compromised. Our research found that senior executives in all organisations included in our study – NFPs, government agencies, and large commercial entities – are very concerned about the risk to data security posed by third-party vendors.

When an organisation provides a third party with the sensitive personal information of their customers or clients, they relinquish control over that data. Organisations attempt to maintain the security of that personal information by requiring vendors to fill out surveys relating to their data protection practices, using various vendor management checklists to assess the security of their vendors. However, it is very difficult for an entity to prove, with any level of certainty, vendors’ compliance, which is often taken at face value, without audit, and without any standards for adherence. This is hugely problematic and carries significant risk. The OAIC’s most recent report from the NDBS highlights the worrying increase in third-party breaches. In the July–December 2023 period, these kinds of breaches increased by more than 400% compared with the previous 6-month period. The OAIC’s report states that:

Most of these multi-party breaches involved a data breach of a cloud or software provider, which then impacted the clients who had outsourced their personal information handling to those providers. This highlights the significant data breach risks that can arise from outsourcing personal information handling. (OAIC, 2024d, p. 28)

Several of our research participants held senior executive positions within organisations that had encountered data breaches due to third-party vendors. In one case, highly sensitive personal health information was discovered on the dark web, which had been compromised because of a third-party breach. The violation of the privacy of the individuals impacted by the breach, and the associated psychological and emotional harm could not be remediated, and the vendor went into voluntary administration because they could not withstand the reputational damage. This scenario demonstrates the negative consequences that can result from third-party breaches.

Referring to the challenge of managing the security of third-party vendors, one interviewee said:

We’re all in the industry struggling with how we deal with third party [risk]. How do we manage that risk in the ecosystem? Breaches often come through the supply chain and the vendors and the third parties that are involved... (Participant 8)

"We're all in the industry *struggling* with how we deal with third party [risk]. How do we *manage* that risk in the ecosystem? Breaches often come through the supply chain and the vendors and the third parties that are involved..."

(Participant 8)



One of our interviewees has worked in information security in Australia for 20 years and, prior to that, in the United Kingdom, and is now the director of a cyber security consultancy that helps organisations increase their cyber security maturity. This participant told us that third-party risk management is particularly challenging, not least because the organisations that engage the consultancy for assistance often do not know how many vendors they have or who those vendors are:

It's something that almost everyone invariably is not doing well enough at. It's something that people have outsourced so much in the past decade, they're struggling to catch up with that. The attackers have absolutely cottoned-on to the benefits of attacking the supply chain rather than the end users ... Vendor risk is hard ... it starts with a bit of a triage process of looking at who your vendors are. Now, a lot of companies cannot even give you a list of who their suppliers are. What they try and do is they look at what bills they're paying or invoices they're paying, and reverse engineer. Then you have to work out what you are paying that supplier to do for you. (Participant 46)

That large organisations are losing track of what vendors they have in their supply chain is very concerning. Engaging third parties significantly increases the risk of data breaches, but organisations can mitigate some of this risk by monitoring their vendors. Organisations losing track of their supply chain to the extent that they are unable to provide a definitive and comprehensive list of their vendors is a serious failing that leaves all stakeholders vulnerable.

One issue relating to third-party risk is the lack of clarity around notification when breaches occur. Individuals' personal information is collected by one entity (the 'data controller') for one purpose and then given to another entity (the 'data processor') for handling. When a breach occurs, it is often unclear as to who is responsible, both for the breach and the associated notification, and sometimes entities will attempt to shift the responsibility for notification and remediation onto the other party. Many interview participants described this problem. One of our interviewees, a senior executive from one of Australia's big four banks, told us that the bank felt like it was being pressured to take responsibility for another firm's breach.

The OAIC's recent report, released on 22 February 2024, speaks directly to this challenge:

In this reporting period, multi-party breaches involving contracted service providers highlighted ... the lack of clearly defined responsibilities should a data breach occur, including who should assess and/or notify the breach ... Prior to using the services of third-party providers ... Entities should ensure service agreements or contractual arrangements address data breach response requirements, including assigning roles and responsibilities for managing a data breach and meeting regulatory reporting obligations. This should specifically address which entity is to assess a data breach should one occur, and which entity is responsible for notifying affected individuals (OAIC, 2024d, p. 29)

This is a problem faced by both entities – the primary organisation and the third-party vendor. Both are impacted in the case of a breach. We interviewed one senior data security professional working for a third-party vendor (or 'data processor'), who expressed difficulty in reconciling the fact that they must rely on the primary organisation (the data collectors and controllers) to consistently maintain adequate privacy, consent, and data collection practices – something that is not always achieved. The interviewee told us:

One of the challenges I face is the fact that for each partner, we agree to deliver everything according to their policies. The consents and the like are agreed [to] with [the primary organisation] as to what the consent will say, etcetera. The privacy policies, the privacy practices and the like that we need to adhere to are [the primary organisation's] privacy policies and practices in each case. It is effectively their consent and privacy position that we are promoting. (Participant 29)

All interviewees said that managing third-party risk to prevent breaches was one of the most challenging parts of information security and this is likely to become increasingly difficult over the next decade, given that entities continue to outsource to an increasing number of third parties. These risks include very significant ambiguities around responsibilities to the regulator and individuals in the case of a notifiable breach. As a result, some breaches go unnotified and for others there is a considerable time lag before the parties agree on a notification approach. The OAIC has signalled third-party risk management as one of its regulatory priorities for 2024.



Recommendation Five:

More rigorous monitoring of third-party suppliers, beyond vendor management checklists, is required. Our research participants commonly said that a consistent standard for security compliance, to which all vendors must adhere, would be beneficial. It is important that vendors can demonstrate that they have the security measures in place through appropriate audit and assurance practices. In addition, there needs to be contractually defined responsibilities for data breach notifications should a breach occur.

In addition, Europe's GDPR distinguishes between 'data controllers' and 'data processors' and creates more clarity around the responsibilities of each in preventing, and responding to, breaches. Following the model set out by the GDPR would enhance Australia's current approach to breaches and breach notifications. Currently, the *Privacy Act* makes no distinction between data controllers and data processors, and views both as APP entities with the same responsibilities and obligations around data notification. However, the GDPR specifies that the data controller (or the primary organisation with which the individual has a direct relationship) is responsible for notifying both the regulator and the impacted individual. This clarity is essential for creating consistency and transparency with notifications across all industries. With the *Privacy Act* currently under review, it is likely that Australia will adopt a similar distinction between data controllers and data processors, but it is important that implementation of the coming changes is swift, particularly given the very significant increase in third-party breach incidents, as reported by the OAIC in early 2024.

"The OAIC's got *guidance*... but it's all written in that *regulatory* way where they don't want any definitiveness in case they get *criticised*, so still, even with guidance, you're trying to work your way through, what are they actually saying?"

(Participant 4)

Data retention: How long to store data?

All interviewees emphasised that data retention poses a significant challenge, increasing the likelihood of breaches and complicating notification practises. Very often, breaches are the result of vulnerabilities in legacy systems, which store the personal information (data) of customers or clients who have received a service from an organisation in the past. In some cases, organisations are legally required to retain old data if it is in the public interest. For example, the *Telecommunications (Interception and Access) Act 1979* (Cth) requires telecommunication companies to retain data for two years to allow law enforcement to carry out criminal and national security investigations. In the financial services industry, AUSTRAC requires entities – including banks, superannuation providers, and insurers – to retain data for seven years. Some data relating to children must be held for longer periods. However, even when an organisation is *not* legally required to retain data, our research found that there is widespread retention of data because it is costly and time consuming for organisations to properly dispose of or decommission old systems. This, coupled with lack of certainty and clarity around how long organisations are legally required to retain data, has created a situation where organisations are retaining data unnecessarily, greatly increasing the likelihood of breaches. The larger the organisation, and the longer it has existed, the more data it is likely to have stored on legacy systems. If a legacy system breach occurs, impacted individuals receive notification from the organisation, but this can be confusing because these individuals have often not had an active account with, or received a service from, the breached organisation for many years.

Interviewees told us that there is a lack of certainty around what data needs to be retained and for how long. This lack of certainty encourages entities to hold onto data in case they are approached by a regulatory or law enforcement body and asked to produce it. Our research also found that there can be disagreement and conflict between different parts of an organisation in relation to data retention practices. Privacy teams are concerned with upholding the privacy of customers and clients, legal teams are focused on the entity meeting its legal requirements (including being able to produce data if requested), and software development teams rely on access to large banks of data for the testing and development of new products (see section 8). Ultimately, organisations frequently choose to retain data because of these competing internal drivers. Most interviewees spoke of a lack of clarity in relation to compliance and the need to establish a best practice approach to data retention. One interviewee said:

... we're all so confused by our destruction obligations. We're being given all different messages ranging from tech department through to the money laundering requirements through to the Corps Act [Corporations Act 2001] through to ASIC. Then throw in the mix regulators who sometimes just randomly send us letters saying 'we require you to stop deleting records because we might be coming to issue you with a demand for records going back 25 years'. People go 'well, I better just hang onto everything'. We're all so confused by our destruction obligations and we're also scared of deleting stuff in case a regulator asks for it so we're hanging onto all this data ... it makes every single breach far more pervasive and significant than it needs to be. In an organisation like ours where we have over 2,000 legacy systems, we've got a 200-year history, the systems don't speak to each other. They don't come with big red delete buttons. (Participant 4)

Safe and secure disposal of data that is no longer required can be expensive for organisations and requires the dedication of time and resources. The time and money required to properly decommission creates a disincentive for organisations. One of our interviewees is a senior cyber security consultant working for a global cloud provider, who previously worked for large Australian financial institutions, and telecommunication and airline companies. They described the problem to us:

If you think of a company that's been around for 20,30,40 years, whose IT systems are going back to the 80s – as you build new stuff, you've gotta either retire the old stuff or upgrade the old stuff. When you get to the cost–benefit analysis, it's a real struggle to justify the investment to uplift [the security of] all these applications, and all of their development processes and all of these copies, and clean all of this information out – especially when you don't even know what they are – to get rid of this stuff. (Participant 28)

Another interviewee told us:

Often, these companies have it in their contract that says 'yes, we'll delete the data', but that's a pain, and it's hard, and it's expensive. They don't bother unless they're made to. (Participant 15)

Whilst the secure disposal of data requires allocated resources, failing to invest in proper disposal creates the risk of much greater negative consequences for both organisations and individuals. When individuals are informed of data breaches in legacy systems, it creates a lack of confidence. These systems should have been decommissioned once the data was no longer required, and this oversight not only financially impacts the entities involved but also burdens the affected individuals.



Recommendation Six:

After the Optus breach in November 2022, Australia's Attorney-General Mark Dreyfus announced that companies would face significant penalties for inadequately managing the storage and destruction of individuals' personal information. This provision is likely to be included in the upcoming amendment to the *Privacy Act*. To ensure compliance, fines must be substantial enough to compel companies to allocate time and resources to proper data disposal practices. Additionally, more comprehensive auditing is essential to prevent organisations from retaining vast amounts of data merely due to cost and convenience. Furthermore, Australians would benefit from the introduction of the 'right to erasure' (similar to Europe's GDPR), granting individuals the right to be forgotten by organisations and have their personal information permanently erased. Australia would benefit from the introduction of a system that requires organisations to lodge annual notifications to confirm that they have complied with their data retention and destruction obligations, including the proper decommissioning of legacy systems. Such notification could include a report of how many requests for data erasure an organisation has received and how this erasure has been actioned.


Policy makers need to reconsider the costs and benefits of retaining data in the current digital climate, and, having weighed up the significant risk that comes with retaining data against the benefits of retention, develop a new conceptual framework around the purposes of data storage and whether, in the digital age, full records need to be stored for potential investigations or whether summarised metadata would be sufficient.

Overall, for the NDBS to be managed effectively for both the organisation and individuals, retention rules must be clarified and streamlined.

Production data testing: Unnecessarily increasing risks?

Organisations, especially large commercial entities, regularly release new products to consumers in order to remain competitive. An example is a new feature in a banking app that allows consumers to more easily categorise and track their spending habits. To develop new products like this, software and tech development teams carry out internal testing to ensure the new feature or product works as intended. Repeated testing with large databases is required. From a user experience perspective, the ideal way of determining whether a new product works properly is to conduct testing with real data (referred to as production data). However, this practice is highly risky as

personal information is vulnerable when placed in insecure environments. For this reason, it is generally recommended that production data should never be used in testing environments. Organisations should 'mask' or deidentify (e.g. pseudonymise or anonymise) the data (referred to as test data) before using it for testing. Our research found that despite the known risks, the practice of using non-deidentified production data in test environments continues. This practice is not consistent with APP 11.1, which requires entities to take reasonable steps to ensure the security of the personal information held (see OAIC, 2019)



"I think there's a *balance* a lot of organisations are struggling to strike between what's the quickest, easiest, fastest and probably the most *accurate* way to test this stuff conversely with the things that could go wrong. "

(Participant 31)

One of our interviewees is a senior executive at another of Australia's big four financial institutions. This research participant told us that they are concerned about what they view as a 'speed to market' culture within the financial sector, in which risky software development practices – like production data testing – are tolerated because they help commercial entities release new products more quickly and, therefore, remain competitive. The interviewee told us:

We saw a lot of those examples. A lot of system development, the use of real customer data for the purposes of developing and testing systems ... I think there is a cultural thing around speed to market, winning with products in the market first. I think there's a balance a lot of organisations are struggling to strike between what's the quickest, easiest, fastest and probably the most accurate way to test this stuff conversely with the things that could go wrong. There are technical solutions out there [that don't require testing with real data] but they just take time to embed, re-engineer, for the purposes of these scenarios, but I think it starts with a culture or an appetite these organisations have for pace and speed without necessarily always knowing what could go wrong, what is gonna go wrong, what has gone wrong. (Participant 31)

Another interviewee, who had previously worked in senior roles for leading financial institutions, said that production data testing is the foremost problem in the information security and cyber security industry. Referring to entities they worked for in the past, the interviewee told us:

It was very, very common practice to do that ... it's the easiest and simplest way to test your application, to know that it's gonna work, to use real data ... the proliferation of that stuff is probably the number-one problem that would be facing the industry – how do you manage nonproduction data, sensitive nonproduction data, and clean it up, manage it ... it's one thing I wanted to call out because I've seen it so much in every industry, to the point where it's systemic ... they might just go 'oh, well, instead of taking a copy, we'll just feed the same information into both systems'. It's literally live, real information going into systems that are not live and real and have low security. (Participant 28)

Our research revealed that privacy and data security can be compromised when actual personal information is utilised during the training of new employees. Similar to the software development practice of testing with production data, entities opt for using real data because it provides the most accurate representation of how a specific process functions. It also requires more time and effort to 'sanitise' or de-identify data sets before using them. Multiple interviewees described this as a significant area of concern.



Recommendation Seven:

The regulatory authority must diligently ensure that organisations adhere to APP 11 of the Privacy Act. Specifically, the use of personally identifiable information – data that remains unmasked or unsanitised – for testing purposes contradicts APP 11. This principle requires an APP entity holding personal information to take reasonable steps to either destroy the information or ensure its de-identification (see OAIC, 2019). The personal information of customers and clients should not be utilised for training purposes.

Production data testing dramatically increases the risk of data breaches. For this reason, the OAIC should require entities to disclose, as part of their breach notification, the fact that the breach has occurred due to this kind of testing. This would give the OAIC some insight as to whether such breaches could be avoided through more careful use of test data. This would also incentivise organisations to move away from these risky testing practises. As with our other recommendations in this report, enforcing safe organisational data practices – critical for safeguarding individuals' information – requires a regulator with significantly increased capacity and authority.

Not-for-profits: Unique challenges

Our research project sought to examine the issue of privacy and data breaches from the perspective of different organisational types: large commercial entities, state and federal level government agencies, and key NFPs. The protection of sensitive personal information and the prevention of data breaches is of particular importance for NFPs that work with vulnerable people. We interviewed senior executives from both large and small NFPs that work with people experiencing homelessness, domestic, sexual and family violence, addiction, and serious physical and mental health disorders. We also spoke to major aged care and disability support providers. Generally, vulnerable people are considered to need more protection and support compared to the general population. However, one of our most important findings is that NFPs are often at greater risk of experiencing harmful data breaches compared with large commercial sector entities, due to the sensitive nature of the personal information they hold. Furthermore, our research has identified challenges related to notification practices for organisations that serve vulnerable populations, including providers of disability and aged care services.

For organisations that work with vulnerable people, the concept of ‘harm’ is more complex and requires more consideration and greater responsibility on the part of NFPs. Data breaches that would otherwise pose minimal threat can be very harmful when vulnerable people are involved, and include physical, psychological, emotional, and reputational harm. These are often more likely to occur as a result of breaches experienced by NFP groups working with vulnerable people because the personal information held by these organisations often includes family histories and their dependents, creating more data points at risk of compromise. Furthermore, holding health-related information makes organisations much more likely to be targeted by cyber criminals. It is known that threat actors are particularly interested in obtaining health-related data because, unlike credit card details that can easily be changed, health information is fixed and contains more personally identifiable information. This explains in part why several Australian hospitals have been the victim of cyber breaches in recent years, for example, attacks on St Vincents Hospital, Sydney, in December 2023 (see Kolovos, 2023) and the operator of four Melbourne hospitals, Eastern Health, in March 2021 (see Cunningham, 2021). Reports show that the value of health-related data can be as much as \$1000 on the dark web, compared to credit card details, which are worth around \$5 (Forbes, 2022). One of our interviewees told us:

... we know that the threat actors are particularly interested in medical information because it's useful for both fraud and identity theft. It carries value on the dark web more than bank credit card numbers ... The more data points you have about a person, the easier it is to pretend to be them, and knowing some of their medical history is really quite valuable. (Participant 8)

There are innumerable small, community-based organisations covered by the Privacy Act and bound by the NDBS because they offer a health service or they collect data relating to health (OAIC, 2024b). The NDBS covers virtually all organisations working with vulnerable people: disability, aged cared, domestic violence, mental health, at-risk children, addiction support, homelessness and people living with HIV and other illnesses. However, NFPs typically do not have the resources and funding to invest in cyber security like commercial entities. The spokesperson for the Australian Information Security Association, James Turner, has previously said that: *‘Not everyone has the resources of a bank to defend against cyber attacks, but everyone is being attacked’* (Turner, 2016). In practice this means that personal information held by NFPs, which is more sensitive than financial information and where there is greater risk of serious harm, is less protected from breaches.

One of our participants is the manager of a small community based NFP organisation that supports people transitioning back into mainstream society following custodial sentences. Our interviewee told us that the data they hold includes:

Case files, sentencing reports, psychologist and psychiatrist reports ... [with] some of the information we have, potentially if it got out someone could take their own life because of the shame. (Participant 37)

This interviewee explained to us that, because of the extremely sensitive health and personal information their organisation holds, it is regularly encouraged by the Government to do more to protect the information, including investing more in cyber security. In this case, given the Department of Communities and Justice is the organisation’s primary source of funds, our interviewee believed that for the organisation to meet growing cyber security demands the associated costs need to be properly provided for within contracts for services.

Every year a *State of the Sector* report (Charity Research Centre Australia, 2023) is released, which examines key trends and issues in the Australian NFP sector. The 2023 report found that one in five Australian NFPs believe that a cyber attack would devastate their organisation. The report also states that 80% of the organisations included in the study have had no recent cyber security training (pp. 9–10). Our research points to the fact that organisations who work with vulnerable people and collect highly personal health information need greater support and more resources to be able to effectively meet their data security obligations.



Recommendation Eight:

Our research revealed a significant disparity between the data protection capabilities of the for-profit and NFP sectors. NFPs, often constrained by limited budgets, struggle to allocate resources for robust cyber security. To address this, we propose several key actions from the sector. First, NFPs should actively seek additional funding from donors, emphasising the critical role data security plays in their mission. Second, when engaging with third-party vendors, contracts should explicitly include cyber security costs. Third, they should prioritise the protection of sensitive personal information, especially health-related data. Fourth, they should advocate for increased government funding, particularly for organisations working with vulnerable groups. Lastly, comprehensive cyber security training for all NFP employees is essential to prevent breaches and enhance protection. By implementing these measures, NFPs can better safeguard data and mitigate risks. This will also help ensure there is greater symmetry in the notification practices of NFPs and for-profit organisations.

Cyber experts: More than a niche skill?

Our research found a concerning shortage of cyber security professionals with the right skills and experience to protect Australian organisations and individuals from data breaches. As a result, organisations are not well placed to make the assessments required to determine whether a breach is notifiable under the law. Some of the most senior executives we interviewed told us that it was difficult to find people with the right experience and skills to effectively manage the complex data security workload, let alone those with the skills to participate effectively in multidisciplinary teams determining critical response decisions if a data breach was to occur. The views of senior executives we interviewed mirrored the findings of research demonstrating that there is a massive skills shortage in this area (Mason, 2022). Cyber security professionals must stay abreast of a constantly evolving security landscape, and people who were trained 10 or 20 years ago must continually retrain to keep up with developments. We were also told that employees with less experience are successfully demanding higher remuneration because of the skills shortage. One participant told us:

I'm just seeing that over and over again, that the demand is high. We're getting people with less experience for more money ... Every time you replace someone now, it's with less experience, less skills for more money. (Participant 35)

Another interviewee told us that it was hard to retain qualified professionals:

Staff retention in these highly skilled, highly sought after areas is difficult ... there's a low supply. (Participant 2)

A third interviewee told us:

... we really need more people, more focus, more energy, really, to stop these things from happening. (Participant 7)

One of the world's largest member associations for cyber security professionals is the International Information System Security Certification Consortium, known as ISC2, which produces an annual Cyber Workforce Study. In the report from their 2023 study, ISC2 said that the global cyber security workforce shortage was just under 4 million. That is, an additional 4 million skilled professionals are needed to fill the current workforce gap, and this shortfall rose by 12.6% between 2022 and 2023. Furthermore, two thirds

of organisations included in the study (67%) lacked the cyber security staff needed to prevent and troubleshoot security issues. According to the report, an inability to find people with the right skills (44%), struggling to keep people with in-demand skills (42%), and lacking the budget to hire people (41%) are the biggest causes for these skills gaps (see ISC2, 2023a). ISC2 CEO Clar Rosso last year said that: *'... the pressing reality is that we must double this workforce to adequately protect organisations and their critical assets ...'* (ISC2, 2023b). Doubling the amount of skilled professionals requires investment by proactive governments and this will become increasingly important over the next five years. Without an increased number of sufficiently skilled and experienced cyber security professionals, appropriate notification practices risk being deprioritised in favour of organisational and operational tasks.



Recommendation Nine:

In response to the growing demand for cyber security expertise, government investment is needed to upskill existing professionals and train many more. Subsidised schemes that encourage more people to train in the crucial areas of cyber security and data protection would help to address the workplace shortfall over the coming decade. Organisations could also focus more on upskilling their existing employees and expanding internship programs to help new graduates gain valuable workplace experience. This approach would ensure that organisations possess the expertise needed to conduct timely and evidence-based assessments of harm, enabling them to notify affected individuals appropriately. This, in turn, supports the implementation of effective risk mitigation strategies for those impacted. Universities need to expand and promote their cyber security course offerings and consider hiring more academics in this space.

Best practice: Optimisation through industry collaboration?

Responding effectively to both the protection of data and effective communication in the context of data breach events is an increasingly complex challenge that is shared across all types of organisations in every industry. Within this ecosystem, it is essential that organisations can adopt a best practice approach to data breach notifications. Our research found that there is a need for more proactive inter-organisational sharing of experiences and lessons from breaches, which will help to bolster all efforts to protect data and to make more robust and timely notification decisions. Many interviewees explained that they learned a lot from observing the Optus, Medibank, and Latitude breaches of 2022–2023 and how these events played out in the media. As previously discussed, interviewees described a ‘culture of embarrassment’ in the corporate sector (see section 2), in which companies that experience breaches are blamed and shamed. Major corporations may perceive breaches suffered by their rivals as advantageous from a financial standpoint.

Our research found that in some cases there has been active and direct lesson-sharing between organisational leaders. This was appreciated greatly by leaders and described as a positive step forward. We interviewed a senior data professional at a major Australian NFP that works with children. They told us:

We do quite a bit of sharing ... I caught up with the CEO of Optus at that time because they'd just been through their breach, and we were in the middle of ours. We hadn't done all of our communications or anything like that, but we wanted to understand what were the things that we can do. She spent an hour with us talking about her learnings etc, her advice to us. We talked to a number of other organisations as well ... because we're in the [not-for-profit] sector we don't compete ... people are willing to share because we're non-competitive and they're very generous with their time with us. (Participant 43)

Inter-organisational sharing seems to be much more common in the NFP sector, compared to the commercial sector. Another interviewee, who also worked as a senior information security professional at an Australian NFP told us:

We do that a bit in the not-for-profit space where we share ideas. We don't necessarily see each other as competitors, but there is a common goal [we are] working towards ... That the more we can share and get people skills and tools to help those most vulnerable people, the better. (Participant 42)

The ACSC, part of the ASD, recognises that protecting data in our modern, highly digitised world is a collective challenge. The ACSC leads the Government's efforts to increase cyber security and protect data. A recent initiative that is helping to facilitate inter-organisational lesson sharing from breaches, is the ACSC's Cybersecurity Partnership Program, which brings together cyber security leaders from across government, industry, and academia to support and learn from each other, aiming to facilitate the sharing of insights and collaboration in the struggle to protect data. The partnership program gives organisational leaders access to threat intelligence and ‘situational awareness’ and runs resilience-building activities.



Recommendation Ten:

All organisations would benefit from more structured opportunities to share experiences and learnings from breaches and subsequent notifications. The ASD runs the ASCPP, which helps to facilitate the sharing of knowledge and expertise to collectively improve cyber security resilience across Australia. The ASCPP's work must be well resourced and ongoing to ensure best practices are shared across industries and sectors, and that the notification practices of organisations are informed by current best practice.


"We do that a bit in the not-for-profit space where we *share* ideas. We don't necessarily see each other as competitors, but there is a common goal [we are] working towards... That the more we can share and get people, skills and tools to help those most vulnerable people, the *better*."

(Participant 42)

Conclusion

Australian organisations face significant challenges in the struggle to protect the data they hold and there are a range of ways organisations can be better supported to ensure that when a data breach occurs, the process of notifying the regulator and impacted individuals is timely, comprehensive, and supports individuals to feel confident that the risks associated with our digital lives are being managed effectively. By interviewing senior executives and other professionals from a range of different information security and privacy roles across sectors, we were able to create a picture of the challenges organisations encounter when faced with a data breach – particularly as they relate to effective notification of the regulator and the people affected by the breach.

Our study underscores the need for proactive measures to improve data breach notifications. By fostering transparency, refining breach reporting practices, and clarifying responsibilities, Australian organisations can better protect sensitive information and contribute to a more secure digital landscape.



By fostering transparency, refining breach reporting practices, and *clarifying* responsibilities, Australian organisations can better *protect* sensitive information and contribute to a more *secure* digital landscape.

References

- Australian Cyber Security* (2024) 'Paying cyber ransoms still lands organisations in hot water', <https://australiancybersecuritymagazine.com.au/paying-cyber-ransoms-still-lands-organisations-in-hot-water/#:~:text=Paying%20a%20cyber%20ransom%2C%20could,in%20'malicious%20cyber%20activity'>.
- Australian Government Department of Home Affairs (2023) *Australian Cybersecurity Strategy 2023-2030*, <https://www.homeaffairs.gov.au/about-us/our-portfolios/cyber-security/strategy/2023-2030-australian-cyber-security-strategy>
- Charity Research Centre Australia (2023) *State of the Sector*. <https://www.uwa.edu.au/schools/-/media/Centre-for-Public-Value/Resources/230906-State-of-the-Sector-Report.pdf>
- Cunningham, M. (2021, March 29) 'Staff unable to access patient files after Eastern Health cyber attack'. *The Age*, <https://www.theage.com.au/national/victoria/staff-unable-to-access-patient-files-after-eastern-health-cyber-attack-20210329-p57eyj.html>
- Fell, J., Piper, G. & Liddy, M. (2023, March 28) 'This is the most detailed portrait yet of data breaches in Australia'. *ABC News Online*, <https://www.abc.net.au/news/2023-03-28/detailed-portrait-data-breaches-oaic-disclosures/102131586>
- Forbes* (2022) *Healthcare Data: The Perfect Storm*. <https://www.forbes.com/sites/forbestechcouncil/2022/01/14/healthcare-data-the-perfect-storm/?sh=62091b376c88>
- Information Commissioner's Office (2024) *Retention and Destruction of Information*. <https://ico.org.uk/for-organisations/foi-eir-and-access-to-information/freedom-of-information-and-environmental-information-regulations/retention-and-destruction-of-information/#:~:text=You%20should%20keep%20information%20for,transferring%20it%20to%20an%20archive>
- Information Systems Audit and Control Association (2019) 'New study reveals cybercrime may be widely underreported – even when laws mandate disclosure'. <https://www.isaca.org/about-us/newsroom/press-releases/2019/new-study-reveals-cybercrime-may-be-widely-underreported-even-when-laws-mandate-disclosure>
- ISC2 (2023a) *Cyber Workforce Study*. <https://www.isc2.org/Research>
- ISC2 (2023b) *Cybersecurity Workforce Study: Looking Deeper into the Workforce Gap*. <https://www.isc2.org/Insights/2023/11/ISC2-Cybersecurity-Workforce-Study-Looking-Deeper-into-the-Workforce-Gap>
- Karp, P. (2023, 19 January) 'Australia to consider European-style right to be forgotten privacy laws'. *The Guardian*. <https://www.theguardian.com/australia-news/2023/jan/19/right-to-be-forgotten-australia-europe-gdpr-privacy-laws>
- Kolovos, B. (2023, 22 December) 'St Vincent's Health Australia says data stolen in cyber attack'. *The Guardian*. <https://www.theguardian.com/australia-news/2023/dec/22/st-vincent-s-health-australia-hack-cyberattack-data-stolen-hospital-aged-care-what-to-do>
- Mason, M. (2022, 13 September) 'Cyber skills shortage to hit 30,000 in four years', *Australian Financial Review*. <https://www.afr.com/technology/cyber-skills-shortage-to-hit-30-000-in-four-years-20220912-p5bhde>
- McGrath Nicol (2023) *Ransomware: A Cost of Doing Business?* <https://www.mcgrathnicol.com/insight/ransomware-a-cost-of-doing-business>

References cont.

Office of the Australian Information Commissioner (2019) *Chapter 11: APP 11 Security of personal information*.
<https://www.oaic.gov.au/privacy/australian-privacy-principles/australian-privacy-principles-guidelines/chapter-11-app-11-security-of-personal-information>

Office of the Australian Information Commissioner (2023a) *Australian Community Attitudes to Privacy Survey*.
<https://www.oaic.gov.au/engage-with-us/research-and-training-resources/research/australian-community-attitudes-to-privacy-survey/australian-community-attitudes-to-privacy-survey-2023#main-findings>

Office of the Australian Information Commissioner (2024a) *Notifiable Data Breaches*.
<https://www.oaic.gov.au/privacy/notifiable-data-breaches/when-to-report-a-data-breach>

Office of the Australian Information Commissioner (2024b) *Privacy for Not-for-profits, Including Charities*.
<https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/organisations/privacy-for-not-for-profits,-including-charities>

Office of the Australian Information Commissioner (2024c). *What is a Data Breach*.
<https://www.oaic.gov.au/privacy/your-privacy-rights/data-breaches/what-is-a-data-breach>

Office of the Australian Information Commissioner (2024d) *Notifiable Data Breaches Report: July to December 2023*.
<https://www.oaic.gov.au/privacy/notifiable-data-breaches/notifiable-data-breaches-publications/notifiable-data-breaches-report-july-to-december-2023#comparison-of-top-5-sectors>

Office of the Australian Information Commissioner (2024e) *Part 4: Notifiable Data Breaches Scheme*.
<https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/preventing-preparing-for-and-responding-to-data-breaches/data-breach-preparation-and-response/part-4-notifiable-data-breach-ndb-scheme#data-breaches-involving-more-than-one-entity>

SurfShark (2023) *Global Data Breach Statistics*.
<https://surfshark.com/research/data-breach-monitoring>

Tillet and Smith (2023, Nov 22) 'Ban on paying hacker ransoms is inevitable, but not yet: Labor'. *Australian Financial Review*.
<https://www.afr.com/politics/federal/ban-on-paying-hacker-ransoms-is-inevitable-but-not-yet-labor-20231122-p5eltz#:~:text=%E2%80%9CEvery%20time%20a%20ransom%20is,hard%20work%2C%E2%80%9D%20she%20said.>

Turner (2016, April 11) 'Australia is suffering a shortage of world-class cyber security teams'. *Australian Financial Review*.
<https://www.afr.com/technology/australia-is-suffering-a-shortage-of-worldclass-cyber-security-teams-20160331-gnv3p7>

World Economic Forum (2022) *Cybersecurity is an Environmental, Social and Governance Issue. Here's Why*.
<https://www.weforum.org/agenda/2022/03/three-reasons-why-cybersecurity-is-a-critical-component-of-esg>

World Economic Forum (2024) 'What does 2024 have in store for the world of cybersecurity?'
<https://www.weforum.org/agenda/2024/02/what-does-2024-have-in-store-for-the-world-of-cybersecurity/>

List of acronyms

ACSC	Australian Cyber Security Centre
ACSPP	Australian Cyber Security Partnership Program
AHPRA	Australian Health Practitioner Regulation Agency
APP	Australian Privacy Principle
APRA	Australian Prudential Regulation Authority
ASD	Australian Signals Directorate
ASIC	Australian Securities and Investments Commission
ASX	Australian Securities Exchange
ATO	Australian Taxation Office
AUSTRAC	Australian Transaction Reports and Analysis Centre
GDPR	General Data Protection Regulation (Europe)
ICO	Information Commissioner's Office (United Kingdom)
ISACA	Information Systems Audit and Control Association
ISC2	Information System Security Certification Consortium
NDB	Notifiable Data Breach
NDBS	Notifiable Data Breaches Scheme
NFP	Not for profit (entity)
OAIC	Office of the Australian Information Commissioner
SME	Small to medium-sized enterprise

About the Authors



Professor Jane Andrew

Jane Andrew is Professor of Accounting, Governance and Regulation at the University of Sydney Business School and Head of Discipline. Jane is co-chief investigator for an Australian Research Council Discovery Grant project examining organisational data breach disclosures.

jane.andrew@sydney.edu.au



A/Prof Max Baker

Max Baker is Associate Professor of Accounting, Governance and Regulation at the University of Sydney Business School. He is co-chief investigator for an Australian Research Council Discovery Grant project examining organisational data breach disclosures.

max.baker@sydney.edu.au



Dr Penelope Bowyer-Pont

Penelope Bowyer-Pont is a researcher in the discipline of Accounting, Governance and Regulation at the University of Sydney Business School. She has a PhD in Political Sociology from Macquarie University.

penelope.bowyerpont@sydney.edu.au



THE UNIVERSITY OF
SYDNEY

The University of Sydney
Business School

www.doi.org/10.25910/psq3-q365

Vol. 1 No. 1, *Data Breaches in Australian Organisations*.

CRICOS 00026A