Privacy Commissioner
Te Mana Mātāpono Matatapu

Inquiry into Foodstuffs North Island trial use of facial recognition technology

Report by the Privacy Commissioner pursuant to section 17(1)(i) of the Privacy Act 2020

MAY 2025

Table of Contents

1 Executive summary

1. Using live facial recognition to identify people of interest is relatively novel in New Zealand. Until now, the only example of its use has been in gaming venues, as a means of identifying problem gamblers. However, the retail sector is actively considering whether the technology could be a useful tool to help businesses combat the rising problem of retail crime, especially serious offences such as violence towards staff and customers, and higher value theft.

2. To this end, in 2024, Foodstuffs North Island ("FSNI") conducted a six-month trial of live facial recognition in selected supermarkets. It also engaged an independent evaluator (Scarlatti) to assess whether the trial had operated successfully and safely.

3. It is hard to overstate the privacy implications of a technology that, if widely deployed in supermarkets, would capture images and process the faces of millions of New Zealanders going about their daily lives. The risks of overcollection, scope creep, surveillance, misidentification and bias are well documented. Biometric technologies such as facial recognition also capture unchangeable aspects of who we are. It is therefore unsurprising that many people have concerns about live facial recognition that is used to identify people.

4. Nevertheless, we acknowledge that like other biometric technologies, FRT can have benefits for organisations and for individuals if it is used safely – with good privacy design, robust processes, and strong safeguards.

5. As a result, the Office of the Privacy Commissioner (OPC) has engaged with FSNI as they developed their operating model. We proposed the trial as a way of evaluating the impact of the model on repeat retail crime offenders and testing the effectiveness of the privacy safeguards. We established a parallel Inquiry so that we could monitor and learn from FSNI's experience and share those findings with others who have an interest in this technology.

6. This report records our findings.

Findings relating to the trial

7. Our overall finding is that the live FRT operating model deployed by FSNI during the trial complied with the Privacy Act.

8. There are several key features of FSNI's operational model (as it was updated during the trial) that enabled us to come to this conclusion:
   (a) **A clear and limited purpose:** FSNI focused the use of the technology on identifying people who had committed serious harmful behaviour in the stores in the recent past: that is, physical and verbal assault, violent and threatening behaviour and higher value theft. No other uses of the information were permitted.
   (b) **The system was effective to address that purpose:** The independent evaluation, combined with our own interviews with staff in trial stores, provides sufficiently strong evidence to conclude that, on balance, using a live facial recognition system was an effective way to reduce serious repeat offending during the trial period. However, the underlying trial data has limitations (see under 119).
   (c) **Fit for purpose technology:** FSNI chose a technology product that had been proved to work at a high-quality level "in the wild", that is in environments where people are not specifically posing for photographs. The technology had not been trained on a New Zealand population since there is no New Zealand specific training dataset. However, it had been trained on similar groups in Australia (including Māori and Pacific people), which reduced the potential for technical bias in the matching process as much as was reasonably possible at this stage.
   (d) **No use of images for system training:** The software vendor was expressly prohibited from using the images collected to train the technology. We strongly support the development of a New-Zealand training dataset, but this should only be done on a consent basis.
   (e) **Immediate deletion of most images:** Images that did not trigger a match against the store's "watchlist" (the image database of people of interest) were deleted almost instantaneously. As Appendix 5 illustrates, this accounted for the vast majority of images that were captured in stores.
   (f) **Rapid deletion of images where no action taken:** Matches that were not actioned were deleted by midnight on the same day.
   (g) **"Watchlists" were generally of reasonable quality and carefully controlled:** Each store has a separate watchlist. Only specific trained staff

were permitted to add people of interest to the watchlist. There were reasonably clear criteria about the types of offences that would lead to someone being enrolled on a watchlist and at least two staff needed to confirm that the criteria were met. Staff were not permitted to add images of children or young people under 18, elderly people, or people with known mental health conditions. Staff were required to err on the side of caution when making these judgements.

(h) **Retention of watchlist information was limited**: There is a retention period of 2 years for information about principal offenders, and 3 months for accomplices. Limiting the time someone can be on a watchlist helps to ensure the information is still relevant and up to date and reduces ongoing impact on individuals.

(i) **Watchlists are not shared between stores**. All watchlists are store-specific and are not shared with other FSNI stores. This helps to ensure that people are not barred from all the stores that they might need to visit to purchase food and other necessities.

(j) **Accuracy levels were acceptable, once adjusted in response to problems:** At the start of the trial, FSNI had set the confidence level at which the system registered a match at 90%, which they considered was high enough to reduce the risks of misidentification to an acceptable level. However, there were two instances where people suffered harm after they were misidentified, and other safeguards did not work as intended. This showed that the initial operating model as a whole did not yet meet the standard required under principle 8 of the Privacy Act (the accuracy principle). FSNI learned from the misidentification incidents. It instructed staff not to consider intervening unless the match level was at least 92.5%. Other adjustments were also made to watchlist image quality, to rules for alert checking, and to staff training. There were no further harmful incidents after those improvements were made, which suggests that the package of safeguards was operating successfully. Such improvements show the benefits of the ‘test/review/learn/update’ approach that FSNI employed during the trial.

(k) **Alerts were checked by two trained staff:** Before deployment, staff were made aware that the technology was not failsafe. Alerts needed to be generated by more than one camera. The alert then had to be checked by two trained staff members in store, who had clear criteria to work with, but who then had to exercise human judgement about whether and how to intervene, or whether to call Police.

(l) **Reasonable degree of transparency that the FRT trial was operating**: Stores had clear A1/A0 signage at the entrance alerting customers that the trial was operating, with more signs in store. Information was published on the FSNI

website. Further information was also available on request at the customer information desk. Staff involved in the trial were trained to answer questions; other staff were trained to direct questions to the customer information desk.

(m) **No apparent bias or discrimination in how discretion was exercised:** From our sample checking, our Office's compliance team found no apparent bias or discrimination in how the watchlists were compiled, in how alerts were checked, or in how decisions were made about whether to intervene.

(n) **Processes for requests and complaints:** People who considered that they had been misidentified or wrongly enrolled on a watchlist were able to make complaints, and have information corrected or removed if a mistake was found.

(o) **Security processes in place to protect information:** Only authorised people had access to the information on the system, and to the security room in which the equipment was stored. There was no automatic connection between the FRT system and the store's standard incident reporting platform: any information had to be manually transferred, and there were clear criteria for what could be added. All access is automatically logged and is regularly reviewed by the Loss Prevention Manager. FRT alerts can only be received by authorised devices which only operate on the in-store network (not from other locations).

(p) **Good record-keeping about system operation:** For instance, records are created about numbers of matches, numbers of alerts, what happens in response to an alert, including decisions whether to intervene (with reasons), how a person reacted if approached, and whether the intervention prevented or triggered harmful behaviour. This record-keeping enables FSNI to continue to monitor the effectiveness of the system.

(q) **Stores have good security infrastructure and are committed to privacy measures**: The stores using FRT have the required level of CCTV technology and cameras in store and the necessary security arrangements including a dedicated security room to accommodate FRT. They demonstrated an awareness and regard for privacy, and the need to comply with the FRT rules and protocols developed by FSNI, including having a store champion.

9. Following discussion with OPC, FSNI finalised a detailed privacy impact assessment ("PIA") as part of its trial planning, which enabled it to identify key areas of risk and create processes to manage most of the risks successfully.

Further improvements needed

10. While the trial model complied with the Privacy Act overall, our Inquiry identified further improvements that would need to be addressed before FSNI considers using FRT permanently or expanding it into additional supermarkets.

    (a) **Update the match algorithm so an alert is triggered at a higher accuracy level:** The current algorithm is still set to trigger an alert at 90% match accuracy, though staff are instructed to observe only and not to intervene unless the match accuracy is at least 92.5%. This was understandable while the trial was under way: system changes can be costly, and there was little point adjusting the algorithm unless the 92.5% was shown to work. However, the higher match level has proved to be one important protection for people, while still enabling FSNI to use the system effectively to address serious offending. The system should therefore be reset to trigger an alert at a minimum of 92.5%. FSNI should then also review the threshold at which staff must take care before acting on an alert (for example, Scarlatti suggested that a match level of 94% may be preferable before staff intervene in response to an alert).

    (b) **Watchlist criteria should remain consistent with store practice during the trial that targeted genuinely harmful behaviour:** FRT is an inherently invasive tool that should be reserved for serious retail crime behaviour, such as offences involving violence, aggression, intimidation or high value theft ("qualifying offences"). This was the case during the trial. However, looking to the future, we do not consider that it would be appropriate to use FRT to manage lower level criminal behaviour such as minor shoplifting, or against people who are perceived as "difficult". It is therefore important that the language of watchlist criteria reflects the high threshold (rather than using more ambiguous terms) to help to ensure that staff who are adding people to watchlists are clear about what is permissible.

    (c) **Check that trespass notices were issued for qualifying offences:** We acknowledge that stores have a legitimate interest in making sure that people do not breach trespass notices. However, trespass notices can be issued for a variety of reasons, some of which are subjective and discretionary, and not all of which are at a level that would justify the application of FRT. It is therefore important to make sure the trespass notice was issued for a qualifying offence before adding someone to a watchlist.

    (d) **Update the privacy impact assessment to reflect what happens in practice:** The initial PIA was an essential tool to help FSNI identify and manage risks. However, some of those settings were altered because of the trial. Both

Scarlatti's evaluation report and this report have also made some further recommendations. The PIA now needs to be updated so that it accurately reflects what happens in practice. This will enable FSNI to use it as a key reference point to guide any broader deployment of FRT in future.

(e) **Continue to review how FRT is operating:** While, on balance, the live FRT model was an effective way to reduce serious repeat offending, there is no evidence about whether it will continue to be effective or justified in the longer term. Crime trends change, as do the laws designed to address them. FSNI therefore needs to keep the use of FRT under active review.

11. As well as making these improvements, it is essential for FSNI to continue to maintain other techniques for managing safety in stores. FRT only works when the agency has a watchlist of people of interest against whom to match images of people entering the premises. However, harmful behaviour can also occur with newcomers to the store, with first-time offenders, or with people who have only committed lower-level offences. It is essential that stores do not rely on FRT alone or see it as a substitute for engaging with the Police, but maintain a range of ways to deal with problems.

What others can learn from this Inquiry

12. Our Inquiry covers the specific privacy issues as they applied to FSNI's operating model for live FRT during the trial. However, many of the trial findings are relevant to other businesses as well as FSNI itself.

13. This report is not a green light for more general use of FRT. However, we recognise the importance of the issue for many businesses. What we have learned from this trial will enable other businesses to ask themselves the right questions about whether FRT is necessary and appropriate for them and will give them strong guidance about what they would need to do to set FRT up in a privacy-safe way.

14. As with any introduction of a technology or process that uses personal information, each new use of FRT must be considered on its own merits and be carefully justified in a business' individual context. It must assess the privacy impacts on members of their community, as well as making sure that FRT will be effective to help prevent crime in a particular setting. All elements of the FRT operating model – not merely the technology itself – must be well designed. This report shows businesses what those elements are and what they would need to do.

15. To be clear, we would not expect every business to trial the technology as FSNI has done. Sometimes FRT will self-evidently be effective to meet a specific lawful purpose. Sometimes, the business environment will be clearly comparable to FSNI's situation and the findings from this trial will be enough to show that FRT would work for them. However, even if a business does not run a trial as such, we recommend testing its FRT settings pre-deployment and over time so it can make any necessary adjustments (for instance to its watchlists, camera placement or staff training) to ensure the system operates effectively and safely.

16. We will continue to monitor FRT adoption and use to ensure ongoing compliance, and we will comment or provide further guidance if necessary. If we receive complaints or become aware of concerns, we will ask the agency to justify its FRT use settings and can take enforcement action if appropriate.

2 Background to the Inquiry

17. Facial recognition systems have many potential benefits in areas such as preventing and detecting crime, controlling access to restricted premises, and verifying identity in areas such as border control. However, they can also raise significant privacy concerns, including overcollection of personal information, expansion of surveillance, inaccuracy or unfair application (including creating or perpetuating discrimination), and lack of transparency for people that the system is operating or what rights are.

18. Any proposed deployment of facial recognition technology ("FRT") must therefore be carefully justified, and the operating model must be designed in a way that either reduces the privacy risks to an acceptable level or removes them altogether. As this report shows, an "operating model" refers not only to the technology itself and the information that underpins how it works, but also all other aspects of what it takes for the process to work well. This includes policies, communications with affected people, staff training, design of interventions, updating records, and handling complaints.

19. The trial of live facial recognition in selected supermarkets by Foodstuffs North Island ("FSNI") was a novel application of the technology in an essential service environment. It presented an ideal opportunity to learn whether such systems can be designed in a way that successfully manages privacy risks. This was valuable not only to inform FSNI's own decision making about whether and how to deploy the technology, but also to inform others who were considering doing the same.

20. FSNI engaged with us early in their thinking and made significant design decisions that reduced privacy risk as a result of those discussions. However, given the significance of the privacy issues involved, and the potential value of the outcome, we recommended that FSNI should run a trial of FRT and we opened a formal Inquiry under section 17(1)(i) of the Privacy Act 2020 (the Act). The trial started on 8 February and ended on 7 September 2024. The Inquiry began on 4 April 2024.

21. While we will continue to watch how FRT is used, including engaging with FSNI and others, this report **concludes** our Inquiry. It records our findings about the trial, makes recommendations for improvements, and identifies what others who are considering using the technology need to consider.

3 What is a facial recognition system?

Main ways of using facial recognition technology

22. There are two main ways in which FRT is used:
    (a) **Verification** uses one-to-one matching. Commonly, it allows a user to verify their identity by matching their face to their own photo identity document or previously enrolled image. For example, your phone uses facial recognition to verify that you are the authorised user, or an automated gate at an airport verifies that you are the passport holder.
    (b) **Identification** uses one-to-many matching. Commonly, it is used to identify whether a specific person appears in a database of many images. In the context of this report, we refer to this database as "the watchlist".

23. It is the second method – one-to-many matching – which is the subject of this Inquiry, specifically as it involves real-time or 'live' collection.

How facial recognition technology works

24. At its simplest level, FRT involves comparing two images of a person's face to determine if the images match. More specifically:
    (a) It creates a mathematical representation of a person's face based on the shape and positioning of key facial features, including eyes, nose and mouth.
    (b) It stores this mathematical representation in a digital summary called a biometric template.
    (c) The biometric template is then compared with a previously stored template that represents that person (in the case of one-to-one matching) or is compared against a database of different images that have been stored on the system to see whether the person appears in that database (in the case of one-to-many matching).
    (d) The system produces a comparison score that measures how *similar* two face templates are and therefore how likely it is that the person is the same.
    (e) If the newly captured image is sufficiently similar, the FRT system records it as a "match", often including an indication of match accuracy.
    (f) The level of similarity required to determine a match is set by the system operator, who then determines what to do, based on that result.

25. Live facial recognition technology performs these actions nearly instantaneously, for example as a person enters a particular area.

The potential for mistaken matches

26. Two types of errors can occur with FRT: ‘false positives’ (where a person is wrongly assumed to be the same as the person on record) and ‘false negatives’ (where the match fails to identify that the person is the same as the person on record).

27. The probabilistic nature of FRT means that people will inevitably be misidentified from time to time. While the system operator must do their best to reduce the error rate, it is therefore equally important to recognise that the system is not failsafe, and to have good additional safeguards to protect people who are misidentified.

28. The accuracy levels of facial recognition algorithms continue to improve and are regularly monitored by the US National Institute of Standards and Technology (NIST), which is the leading international assessor of FRT and FRT vendor products.

29. In ideal conditions, the best facial recognition algorithms can have very low error rates and minimal differences in accuracy across demographic groups. However, FRT is often not deployed in ideal conditions. This is certainly true of supermarkets, where people are on the move and may be wearing items such as hoodies, hats, masks or glasses.

30. There are several factors that can markedly affect FRT error rates:
    (a) **The quality of images captured, and the quality of images enrolled in the watchlist**: The quality of both can be affected by lighting, exposure, camera placement, motion, having multiple people present, or angle of the head. They can also be affected by whether the individual was aware and actively participating in the process (for instance, by looking directly at the camera), or whether they were going about their normal activities and not actively engaging with the camera. The latter is known as FRT “in the wild”.
    (b) **The facial recognition system used:** The systems on the market differ significantly in their ability to accurately identify matches. They also differ by how well they perform “in the wild”. Choice of technology is therefore still a key factor.
    (c) **Facial occlusions** such as glasses, masks, beards, or clothing that partially covers features can make it harder to create an accurate template. **Changes in appearance over** time also make a difference, including aging or facial injuries.
    (d) **The size and contents of the watchlist** (that is, the number of enrolled templates). A watchlist may either be overinclusive, or underinclusive.
    (e) **Demographic characteristics**, including gender (with higher error rates for women) and skin tone (with higher error rates for people with darker skin

tones). Those error rates are likely to be higher if the system has not been trained for a particular population. Relevantly for this report, there is currently no specific training data set for the New Zealand population. While demographic discrepancies have reduced over time elsewhere in the world, it is unclear that these improvements have affected error rates for Māori and Pacific peoples. Negative effects can also compound, for example where higher error rates coincide with populations that are already more subject to disadvantage, surveillance, profiling, or being labelled as "persons of interest".

31. Errors matter at a human level. False positives can result in a business acting against the wrong person (such as accusing the person of an offence or denying them access to essential services). False negatives, on the other hand, may lead to a misplaced sense of security, particularly if there is an overreliance on the FRT system. High error rates may undermine the value of having a system at all.

32. Finally, it is important to note that even a system with very low error rates will produce a substantial number of misidentifications if it is scanning the faces of many people. For instance, supermarkets have millions of customer visits every year. A small percentage of mistaken matches still equates to many affected human beings. This is why other aspects of the system, such as processes for human checking and decision making are so vital. The technology is only part of the picture.

How FRT is used in New Zealand

33. In the last 15 years, facial recognition for verification has started to become more commonly used in New Zealand across a variety of contexts. For instance, in 2009, the New Zealand Customs Service introduced facial recognition border control gates at airports. In 2017, Apple launched the first iPhone that could be unlocked using facial recognition. FRT is also increasingly used for identity verification in banking and anti-money-laundering processes, access to government services using RealMe, and to create digital identities.

34. The only previous example of live FRT being used in New Zealand for one-to-many identification was in gaming venues, particularly casinos. It was first trialled in 2018, to help venue operators identify problem gamblers who have been excluded. It appears to be reasonably effective, and its use is now routine. This use of live FRT in this context seems to be broadly supported by the government, gaming venue operators and the public. This is because of the serious harm to problem gamblers themselves and their families that FRT is deployed to prevent and the legislative

obligations on gaming venues to minimise harm. Importantly from a privacy perspective, the highly targeted nature of the system and its use in limited venues makes it easier to justify its use. In addition, problem gamblers who choose to participate in the self-exclusion programme can actively use it as a tool to help themselves.

35. While many business operators have expressed interest in the potential of live FRT, it does not appear that any other businesses have adopted it until FSNI ran its trial in selected supermarkets.

36. Notably, the New Zealand Police do not currently use live FRT. Their August 2024 assessment indicated that the overall risks currently outweigh the potential benefits in the policing context. It also indicated that further decisions about Police use of live FRT will not be made until the security, privacy, legal, and ethical perspectives are fully understood, and affected communities are consulted.

37. However, the Police do use *retrospective* facial recognition which involves running facial recognition on an image after the image has been captured. For example, Police may use retrospective facial recognition on CCTV footage to see if they can identify an offender after the incident has been bought to their attention.

FRT and retail crime

The rationale for the use of live FRT

38. Retail crime is a major problem in New Zealand, as it is elsewhere. FSNI and other retail operators report a significant increase in violence and assaults associated with theft in stores. High value theft and organised shoplifting is also a problem. To use Foodstuffs' own numbers, total retail crime incidents at FSNI's stores reached 5,124 in Q1 2024 (Jan to March 2024). Within that, violent and aggressive offences (such as assault, intimidation and harassment) are reported to be double the numbers of the previous quarter. Some incidents result in serious injury. In addition to the physical and emotional impact of these incidents on staff and customers, FSNI reports that these incidents deter owners from continuing operation because of the risk to personal safety of staff, customers and owner operators, and the resulting loss of profit.

39. FRT is seen as a potential tool to help to address retail crime by accurately and efficiently identifying people who have previously been involved in serious retail crime incidents. The aim is to prevent reoffending in stores by deterring previous criminals from entering, by enabling staff or Police to intervene before harm is caused where it

is safe to do so, or by providing evidence to support an investigation where intervention is not safe or possible. Successful deployment of FRT may also deter first-time offenders from committing crime because they understand that, if caught, they would be stopped from entering the store in future.

40. Without FRT, staff rely on the following techniques to manage offenders:
    (a) using CCTV to view and record incidents or suspicious activity in store
    (b) relying on security staff to identify repeat offenders from memory or a collection of photographs
    (c) equipping staff with body cameras with video and/or sound to record incidents when they occur for evidential purposes
    (d) training staff on de-escalation techniques.

41. While these techniques are still valuable, FSNI and other retailers do not consider that they are effective enough to deal with the types and scale of retail crime that businesses face. In contrast, FRT provides a proactive, quick and (usually) accurate identification at the point the person enters the store, or shortly afterwards. This gives staff additional time to decide whether and how to act. To illustrate this, FSNI calculated that using the FRT system would give staff an average of 4 extra minutes of time to decide how to respond, compared to a manual recognition system whereby staff have to identify people from memory. It also provides more time for the Police to be called to attend if the situation warrants it.

42. The serious problem of retail crime and the techniques available to manage that problem are highly relevant to this Inquiry. First, they are relevant to how the privacy principles apply to the use of FRT in supermarkets. Also, under section 21 of the Privacy Act, we are required to consider privacy interests alongside other human rights and interests, including how businesses can achieve their objectives efficiently.

Limits of an FRT system in a retail context

43. However, it is important to remember that FRT is not a silver bullet. It will not completely solve the problem of retail crime.

44. First, however good the technology itself may be, an FRT system may fail if other elements of the operating system (such as watchlist creation, alert checking, intervention measures and training on how to use the system) do not work properly.

45. Secondly, it also does not replace other store crime and harm prevention measures. Beyond a possible deterrent effect, the technology alone does not prevent or reduce

incidents of violence, harassment or shoplifting. Its ability to address retail crime is dependent on its deployment as part of an overall security strategy – that is, whether and how staff respond if the system triggers an alert. It is also still primarily the role of the Police as a well trained and equipped organisation with specific statutory powers, to reduce and respond to harmful incidents in stores.

46. There are other limits on how FRT can help retailers to prevent or respond to harmful incidents in their stores:
    (a) FRT will only identify people who are already on the store's watchlist. In the context of this trial, that means that FRT will only identify repeat offenders who have previously engaged in serious harmful behaviour that has warranted them being put on the watchlist. The system will not identify first time offenders. Nor will it identify repeat offenders who were not able to be enrolled in the watchlist (for example because a CCTV image was of too low a quality to create a viable biometric template).
    (b) Once staff are alerted to a trespassed or violent individual, the aim is for them to take an action that will *avoid* further harmful behaviour occurring. However, there is always the risk that the action that staff take (such as asking a person to leave) could cause or escalate an incident. In this situation, the availability of Police to attend the incident can be critical.

47. Retailers also need to be aware that FRT systems operate on the assumption that people who have previously engaged in harmful behaviour or who have committed serious shoplifting are so likely to repeat their actions that the situation warrants them being asked to leave or at least to be monitored as they shop (a form of profiling). While many people may indeed continue to offend, previous behaviour is not necessarily determinative of future behaviour. This is not simply a privacy or ethical problem: individual stores need to make their own choices about whether use of FRT is right for their communities. For some, there will not be sufficient justification to use it.

48. Many retailers may also not find that FRT is a useful investment. Particularly in smaller communities, staff may already instantly recognise repeat offenders who come into their stores, and FRT may not add to either the deterrent or investigative value that existing CCTV provides.

Testing whether FRT helps to address serious crime in supermarkets

49. Theory is one thing: real life conditions are another. This trial has provided an important opportunity to test how effective FRT may be in the context of supermarkets. In particular, it focused on whether developing an FRT operating system may be an effective additional tool to help reduce serious and harmful criminal behaviours by repeat offenders (that is, high value theft; burglary or robbery; verbal or physical assault; and aggressive, violent and or threatening behaviour to staff and customers).

The significance of the supermarket context

50. The supermarket context in New Zealand is particularly unique, with only three major grocery retailers operating here. The two Foodstuffs cooperatives operate only in the North and South Island respectively which means that each main island of New Zealand has only two major grocery retailers operating – a Foodstuffs co-operative (including New World, PAK'nSAVE and Four Square brands) and Woolworths New Zealand (including Woolworths, Fresh Choice and SuperValue brands).

51. As the COVID-19 experience demonstrated, supermarkets are essential services, providing food and other products that are vital to the welfare of a population. In modern societies, people are heavily reliant on supermarkets to obtain their food, and the vast majority of people visit a supermarket regularly. Foodstuffs South Island reports that it serves over 600,000 customers every week, and Foodstuffs North Island reports that it serves an average of 4.08 million customers every week, including online customers. Commerce Commission research shows that, in a typical week, 95% of New Zealand consumers will shop at a supermarket owned by one of the major grocery retailers (Foodstuffs North Island, Foodstuffs South Island or Woolworths NZ). Fewer than 0.5% of shoppers will only visit other stores.

52. In other words, in a typical week, a very large proportion of the New Zealand population will enter a supermarket owned by either a Foodstuffs cooperative or Woolworths NZ. What this means is that in some areas of the country, people who are barred from shopping at their local supermarket may have few other choices about where to shop or what they can buy, at least without travelling long distances or paying for delivery services. Their access to essential services may be restricted.

4 Privacy and social concerns associated with live FRT

53. FRT systems have a variety of applications, and some are less privacy intrusive than others. However, at the more intrusive end of the spectrum, use of facial recognition can have direct and substantial impacts on the rights and interests of members of the public.

54. Use of the technology must therefore be clearly justified, and its operation must be controlled in a way that successfully manages the privacy risks.

Broader social concerns

Concerns related to surveillance

55. Live facial recognition can be particularly intrusive as it can be used to monitor large groups of people in real time. There are concerns that widespread use of FRT could normalise a culture of surveillance, where individuals are increasingly being watched by others and are aware that they may be watched. This has significant consequences such as exacerbation of existing power imbalances and lack of control, especially for people in vulnerable populations. Increased surveillance can have a chilling effect on the exercise of human rights, particularly freedom of movement, association and expression, if people feel that they may be identified and profiled as they go about their daily lives.

56. FRT can also operate without people knowing that it is there. Unlike use of a biometric technology such as fingerprint scanning, FRT operates from a distance, collecting face images remotely via a camera or from a CCTV image. The person does not need to physically interact with any hardware. FRT enables people to be monitored in a way that removes any ability to exert choice or control over the collection and use of their personal information.

57. While many FRT systems function in public or quasi-public spaces, this does not mean that privacy concerns are irrelevant. Use of live FRT for identification purposes is different from simply being able to notice a person who is passing on the street, or even from using CCTV to monitor people in public places. It involves a deliberate search for people of interest to make formal decisions about them, to add to information that is already known about them, or to make assumptions about them (profiling).

58. In the context of supermarkets, use of FRT means that the system will capture an image of every person entering the store, and potentially multiple images of them while they are in specific areas of the store. What the operating system then does with the image or images is critically important to identifying whether the level of privacy risk is high or low. How correct is the match, or the other information associated with that person? How will it be used? How long is it kept for and why? Who is it shared with? Are people told that the systems are operating? Are there unfair negative impacts on vulnerable populations such as children, the elderly, disabled people or others? Are there unfair negative impacts on different cultural groups? Is the system constantly reviewed to make sure it is still operating as intended?

59. The answers to all these questions will help to determine whether the tool is operating in a justified and proportionate way, or whether its use amounts to more intrusive surveillance.

Understanding the impact on Māori

60. Under section 21(c) of the Privacy Act, we are required to take account of cultural perspectives on privacy. We are particularly aware that Māori have significant concerns about the use of FRT.

Risks of inaccurate matches

61. A key concern for OPC at the start of the Inquiry was the impact of the known deficit with FRT globally regarding the accurate identification of people with darker skin tones. The technology used in FRT relies on light reflectance from skin and, while the technology is continually improving, it is still known to be less accurate when responding to light reflecting from darker toned skin.

62. While the accuracy of many systems is improving, there is currently no New Zealand population dataset on which to train and improve FRT systems. As a result, we cannot be completely confident that the technology itself has addressed this inherent deficit sufficiently to prevent harm to New Zealanders, including many Māori and Pacific people.

There are already high levels of concern about impact of the technology

63. OPC's 2025 survey of awareness, knowledge and levels of concern regarding privacy showed that Māori are more likely to express concern regarding facial recognition than other New Zealanders. The results also showed a higher level of concern amongst Māori regarding the potential for technical bias inherent in facial recognition leading to misidentification (for example, being less accurate for darker skinned

women). The same survey found that Māori were more likely to avoid visiting a particular place due to surveillance concerns.

64. These concerns were also raised by Māori stakeholders during consultation on the proposed biometrics code of practice and in public feedback on the FRT trial. In particular:

    (a) There are tikanga Māori considerations with the taking and storage of biometrics. Biometric information has significance and sensitivity for Māori because it is connected to parts of the body or images that are tapu, including tā moko, mataora or moko kauae. This information connects individuals with their whakapapa and ties them to whānau, hapū and iwi, and its use is restricted. Capturing biometric information without consent is felt to be a breach of that person's tino rangatiratanga, tapu, mana and mauri.

    (b) Māori are already over-monitored. As a marginalised and youthful population that is often categorised by race, Māori also carry higher risks of data harms and misuse, which are amplified when biometric technologies are involved. Such technologies can have a disproportionate impact on Māori and Pasifika people, because they work with data that is already biased. Technology can simply make existing bias more efficient.

    (c) They are particularly concerned that these harms will not only be experienced by individuals, but also by groups. Again, existing biases often result in labelling people as ‘problematic’ – a problem that disproportionately affects Māori.

The views of the Māori Reference Panel

65. The Commissioner established a Māori Reference Panel at the end of 2024 to provide OPC with advice on Māori perspectives on privacy issues. This is one way (though not the only way) that the Office uses to enable it to take account of cultural perspectives on privacy, as required by section 21(c) of the Privacy Act. In March 2025, OPC asked the Panel for its views about FRT in supermarkets.

66. As a point of principle, the Panel did not support use of FRT in supermarkets, given the vital role of supermarkets in providing access to food, the current supermarket duopoly which means there are limited alternative options for people who are barred from entry, and the concern that the whole population of Aotearoa will be subjected to surveillance in supermarkets in order to reduce instances of harmful behaviour by a small minority of customers. Concern was also expressed about FRT potentially being used by retailers as a measure to protect revenue rather than safety.

67. They also re-emphasised the following points:

(a) There are risks that stereotypes and resulting human bias will mean Māori will be over-represented on watchlists.

(b) Good data is needed to verify whether the use of FRT is accurate and unbiased across different skin tones. A specific concern was the need for data about bias and ethnicity, or at least data about outcomes for people other than Caucasian people.

(c) There is a need for ongoing staff training on how to verify matches.

(d) There are risks of discrimination and cultural appropriation from capturing mataora and moko kauae, and particular risks of enrolling and targeting rangatahi.

(e) It is important for people to have access to good information about the training data used by FRT providers, given the potential for misidentification for people of colour.

(f) Processes for addressing the harm caused by misidentification of customers should reflect tikanga Māori to restore relationships and balance, for example hohou te rongo.

68. We value the advice that Māori have provided during this process. Their advice has been important to help to ensure that our expectations of whether and how to implement FRT are appropriate.

How the Privacy Act applies to FRT

Biometric information is sensitive

69. An image of a person's face which is collected by FRT is clearly information about that person. Moreover, it can be used to verify or identify that person: indeed, that is the purpose of capturing the image. FRT images are therefore clearly "personal information" and are regulated by the Privacy Act 2020.

70. Unlike some of our overseas counterparts (such as Australia, the EU and the UK), the New Zealand Privacy Act does not define specific categories of personal information as "sensitive". Instead, the information privacy principles (IPPs) create higher expectations for how more sensitive information is treated. The more harm that can be done by mishandling that information, the stronger the protections need to be.

71. The Privacy Commissioner has consistently treated biometric information as sensitive personal information. This is because biometric information is based on the human body and is fundamental to an individual's personhood: it makes you, "you". It differs from many other types of personal information in that it relates to innate

characteristics of a person which are not consciously determined and cannot be easily changed (if they can be changed at all). As such, biometric information requires particular care and respect under the Privacy Act, recognising the inherent dignity of all people.

General principles

72. The use of FRT systems is governed by the Privacy Act 2020 with regulatory oversight by the Privacy Commissioner. Other legislation may also apply in the context of government use of biometric information but is not relevant to the deployment of FRT in supermarkets.

73. The purpose of the Privacy Act is to promote and protect individual privacy. At the core of the Act are the 13 IPPs. These set out the obligations of agencies that handle personal information, and the rights of people who are the subjects of that information.

74. The IPPs are technology neutral. This enables them to apply to a wide variety of different and emerging technologies that collect, match, create, use or rely on personal information, including FRT.

75. In October 2021, OPC published a position paper for public and private sector agencies using or proposing to use biometric systems and biometric information and the relevant considerations under the Privacy Act. The position paper is a useful resource for those considering implementing FRT, as well as this report.

76. In future, FRT and other systems that process biometric information are intended to be governed by a specific Code of Practice under the Privacy Act, issued by the Privacy Commissioner. The Code will modify how some of the IPPs regulate processing of biometric information. However, at time of writing, that Code is not yet part of the law. The law that applies to this Inquiry is therefore the Privacy Act 2020, as it stands.

77. A quick summary of key obligations under the IPPs is as follows:
    (a) Agencies must only collect information that is genuinely necessary for their specific lawful purposes (that is, using the information will be effective to meet that purpose and they cannot deploy a less intrusive alternative) **(IPP1)**.
    (b) Collection must be lawful, fair, and not unreasonably intrusive **(IPP4)**.

(c) When biometric information is collected directly from individuals, they must be informed about how and why it is being collected, who will have access to it, and how it will be stored **(IPP3)**.

(d) Information must be transmitted and held securely to protect it against loss, unauthorised access, and other forms of misuse (including when the agency entrusts the information to a third party to process on its behalf) **(IPP5)**.

(e) Agencies adopting FRT need to ensure that the information that is held, generated and used is sufficiently accurate for the specific purpose for which FRT is being deployed. This applies throughout the operating system: from ensuring the watchlist contains only relevant and accurate information, to accuracy of the match and any subsequent alerts, and actions that are taken in response to those alerts (such as locating the right person in a store) **(IPP8)**.

(f) Information should not be retained for longer than is necessary for its core purpose **(IPP9)**.

(g) Biometric information should not be used or disclosed for different purposes than the core purpose, unless a valid exception in the IPPs applies **(IPP10 and 11)** and should not be disclosed outside NZ unless an exception applies **(IPP12)**.

78. Agencies also need to have a breach management system so they are well placed to meet their obligations to respond to serious privacy breaches and report them to OPC.

79. To meet their obligations in the IPPs, agencies must also take a series of key steps, including:

(a) comparing potential technology solutions to find one that will be fit for the purpose for which it will be used, including being effective, accurate and minimising the risk of bias

(b) conducting any necessary due diligence to ensure that the information will be handled appropriately

(c) adequate testing of the technology prior to adoption to make sure it performs as expected and that any problems can be addressed before it is deployed into a real life situation

(d) carrying out a Privacy Impact Assessment so that privacy risks can be identified and fully managed

(e) monitoring results and performance on an ongoing basis to ensure that the technology is still operating as intended and so that privacy adjustments and further risk mitigations can be incorporated

(f) addressing any complaints that are made about the operation of the system from affected people, including correcting information that is inaccurate so that mistakes cannot recur, supporting them if their information is lost or misused, adjusting the system to prevent mistakes from happening again, and providing any appropriate redress where the person has suffered harm.

80. Often, consultation with relevant stakeholders and impacted groups will also be necessary to identify relevant impacts on individuals and to ensure that privacy safeguards are appropriate to address disproportionate impacts. Stakeholders may well be able to point out relevant issues that help the agency see the full context.

81. Getting the privacy safeguards right is critical to gaining the benefits of FRT while avoiding unintended consequences and ensuring it does not intrude on peoples' privacy. The OSAC technical guidance for implementing FRT explains how clear settings and privacy safeguards will help to limit the privacy intrusive impacts of the technology:

> *…it is not possible to identify people who walk past the system if they are not on the watchlist. Images of subjects who pass the system are not collected for additional analysis and are not added to a watchlist. Nor is it possible to track people as they go about their daily lives. The footprint of the technology should be relative to a specific operational use and location. Whilst passive live facial recognition (LFR) requires automated processing, decisions with regards to identity confirmation are made by the human reviewer and not by the LFR system. Only a very small percentage of people who walk towards an LFR system camera will generate an alert and not all of those alerts will result in an engagement.*

The international context on data protection issues with biometrics

82. New Zealand's privacy framework for FRT is broadly consistent with other comparable countries, although there are some differences in how privacy laws apply. Privacy regulators in other countries have also been examining the private sector use of FRT and have identified areas of concern or non-compliance which New Zealand has been able to learn from. Examples include:

(a) the British Columbian Information and Privacy Commissioner's investigation of Tire Associate Dealers, which found there was a lack of a reasonable purpose for using FRT, and found deficiencies in consent and notification

(b) the UK Information Commissioner's investigation of Facewatch, which identified areas of concern that were addressed through, amongst other things, tightening

when FRT could be used, including focusing on repeat offenders and significant offences

(c) the US Federal Trade Commission's investigation of Rite Aid, which prohibited use of FRT for surveillance purposes

(d) the Australian Privacy Commissioner's determination about the use of FRT by Bunnings, which concluded that Bunnings had not met the threshold for collecting biometric information without consent (either on safety or law enforcement grounds).

83. Despite differences in the specific FRT regulatory frameworks, these examples provide useful benchmarks for adopting FRT in a privacy consistent manner.

Proportionality

84. It is becoming increasingly common for data protection laws (such as Europe's General Data Protection Regulation) and privacy regulators (such as the Australian Privacy Commissioner) to consider whether collections of personal information and its use in technology systems are proportionate to the benefit that the agency is aiming to achieve. For example, an agency's purpose that only achieves a limited benefit but would have a significant impact on peoples' privacy, may be a disproportionate use of a biometric technology such as FRT.

85. While not expressly articulated in New Zealand's Privacy Act, proportionality is a thread that runs through the structure of the IPPs. As well as thinking about effectiveness and intrusiveness, OPC also advises agencies to think about the broader social and systemic impacts of their actions. For agencies considering using FRT in future, it is worth noting that OPC's proposed new Biometrics Code includes an expanded legal obligation to consider broader aspects of proportionality.

86. Thinking about proportionality is particularly important in areas such as FRT that can be perceived as – or even used for – intrusive or unjustified surveillance. Some activities may be so inherently harmful that their benefits are not worth the risks to society or to particular groups within the community. In the context of FRT, we would advise agencies to consider broader proportionality questions as part of their Privacy Impact Assessments, to help them identify and mitigate any unintended or disproportionate consequences.

5 The trial and the Inquiry

Deciding to run a trial

87. OPC was approached by FSNI in June 2021 to discuss the rollout of FRT in their North Island stores. During this early engagement, we discussed:
    (a) the draft privacy impact assessment
    (b) whether the identified problem that the use of FRT was intended to address had been clearly identified, specifically what behaviour the FRT was intended to address
    (c) whether other options had been considered and why they were not adequate to address the problem
    (d) whether FRT would be effective at addressing the identified problem
    (e) the process around the person of interest database ("watchlist"), including how stores decide when to enter people on the register, how to appeal a decision to be put on the register and how long people would stay on the register, and
    (f) whether appropriate stakeholder consultation had been carried out, particularly with Māori and people in the community.

88. In response to OPC's feedback, FSNI made several changes to their PIA and adopted a variety of privacy mitigations.

89. We shared FSNI's concerns about safety in supermarkets, and appreciated why there is a need to find new ways to reduce the harm that was happening. However, we remained concerned that there was a lack of evidence that FRT would be an effective way to address those concerns. The fact that this would be the first use of live FRT outside of problem gambling and in an essential service industry compounded our concerns. We therefore recommended that if FSNI wanted to proceed with the introduction of FRT, it should be introduced initially in specific stores on a time-limited trial basis. Additionally, FSNI should be prepared to stop the use of FRT if the trial showed that it was non-compliant, ineffective or disproportionate when implemented.

90. FSNI agreed. It paused the rollout of FRT to North Island stores generally and agreed to first implement FRT in some stores on a trial basis. The aim of the trial was expressly to test the effectiveness of FRT and the adequacy of the implemented privacy protections. FSNI also halted the use of FRT in some stores that had already been using it, pending the outcome of the trial. Some of these stores later became part of the FRT trial.

Privacy-related changes that were made pre-trial

91. Changes that were made to FSNI's original settings following our discussions included:
    (a) introducing automatic and immediate deletion of biometric information where there is no match with an individual on the watchlist (the initial retention policy was to keep non-matched images for five days)
    (b) narrowing the definition of 'harmful behaviour' that the use of FRT was intended to address, to exclude behaviour that is merely disruptive
    (c) limiting the scope of individuals other than the primary offenders who could be included on a watchlist, by renaming the category of 'associates' to 'accomplices' and restricting it to people who actively assisted in harmful behaviour
    (d) developing criteria and processes for stores to use when responding to a request to be removed from the watchlist
    (e) undertaking further engagement with FSNI staff and First Union
    (f) talking to the Indigenous Genomics Institute about concerns from a Māori perspective, and establishing a Māori Advisory Group for the trial with some consultation with iwi
    (g) reducing the number of stores that would be part of the trial from 32 to 25
    (h) agreeing to employ an independent evaluator to design and assess the trial.

Selection of stores

92. FSNI selected 25 stores from across the co-operative to participate in the trial. All stores operated under the New World or PAK'nSAVE brands.

93. Under FSNI's franchise operating structure, each store is responsible for meeting its Privacy Act obligations, including ensuring that its use of FRT complies with the Act. FSNI's Support Centre provides expert privacy assistance to stores with more complex queries, as well as advising on complaints and IPP6 requests. In the context of this trial, FSNI led the development and implementation of the FRT operating model with relevant documentation (FRT User Manual, training and trial monitoring) provided to stores by FSNI. This central process was important for consistency and avoided duplication of effort.

94. The process for selecting stores considered:
    (a) Number and nature of incidents of harmful behaviour
    (b) store willingness to participate and likelihood of complying with the rules and protocols set out by Foodstuffs, including the quality of reporting required

(c) the store's awareness and regard for privacy
(d) the ability to support the trial by providing a store champion
(e) if it had the required level of CCTV camera system and technology in store
(f) the security arrangements including the security system and if it had a dedicated security room to accommodate the separate FRT system.

95. Our own discussions with store owners also found that stores opted in for consideration for selection based on the number and nature of the incidents occurring and the resulting significant concerns store owners had for staff and customer safety and wellbeing.

Inquiry terms of reference and timeline

96. Given the importance and potential implications of the FRT trial, we opened a formal Inquiry at the time the trial started. Its purpose was to gather information to monitor the implementation of the trial to ensure it complied with the Act, and to ensure that privacy risks were identified and mitigated as set out in FSNI's privacy impact assessment. A key focus of the Inquiry was to see whether the trial provided evidence that suggested FRT is an effective tool in reducing serious incidents caused by repeat offenders in stores.

97. FSNI provided information to this Inquiry under sections 87 and 203 of the Privacy Act. The interviews with store staff were conducted on a voluntary basis. All information provided by FSNI is subject to section 206 of the Privacy Act that requires OPC staff to maintain the secrecy of that information.

98. The Terms of Reference for the Inquiry are set out in Appendix 1, and a detailed timeline of the trial and the Inquiry is set out in Appendix 2.

Inquiry methodology

99. We conducted the Inquiry based on:
(a) documentation and statistics provided by FSNI and its independent evaluator, Scarlatti
(b) site visits, including viewing the FRT system in operation and interviewing key staff
(c) public feedback including complaints received by FSNI and public feedback received by OPC.

Documents reviewed

100. We reviewed the following documents, provided by FSNI:

(a) the operational protocols and the facial recognition user manual

(b) the Privacy Impact Assessment

(c) the findings of Scarlatti as FSNI's independent trial evaluator, at key points during the trial and at the completion of the evaluation report

(d) written feedback from customers or staff, or complaints received by us, by trial stores or by FSNI

(e) feedback from Māori communities to identify any impacts particular to them.

Statistical information that we reviewed

101. We reviewed the following statistical information from the FRT trial stores:

(a) The total number of faces scanned by FRT during the trial period

(b) Total number of facial images deleted including non-matched images

(c) Numbers enrolled on watchlist

(d) Rationale for enrolment by specific category

(e) Action taken in response to identification.

102. Other data and information that was provided to us during our Inquiry included:

(a) Confirmation of the location and number of non-FRT stores

(b) The security measures in place to address harmful behaviour at the non-FRT stores and how it works

(c) The total number of incidents arising from harmful behaviour reported by stores including to the incident reporting platform or by other crime prevention tools such as body-worn cameras or CCTV

(d) Number of customers visiting non-FRT stores during the trial period and the number of serious harm incidents reported in those stores

(e) Number of serious harm incidents broken down by the categories assigned by the evaluator

(f) Evidence of the numerical data for calculating the results that indicated FRT reduced harmful behaviour

(g) Weekly and/or fortnightly reports on FRT trial stores' data and progress

(h) Evidence of their statistical model outputs run to calculate the confidence interval for their report.

103. For the non-FRT stores, we requested comparative information including estimates of total number of customers, and comparative information from the stores' standard incident reporting platform.

Site visits

104. OPC Compliance Officers visited 10 FSNI stores that participated in the trial. We selected these stores with a view to assessing how FRT performed in diverse communities, including where demographic information suggested there were many people who could be disproportionally impacted by the FRT. We also selected communities with higher volumes of reported incidents, on the basis that this would provide more opportunities for the full FRT system and operational processes to be reviewed against documented processes and tested through demonstration.

105. Each visit consisted of an interview with key staff; a demonstration of the FRT system and a discussion on their experiences with the use and effectiveness of FRT (in comparison with other store security techniques) as well as sharing stories of the nature of serious harm incidents occurring in each store.

106. Demonstrations of the FRT system included reviewing placement of cameras and reviewing the contents of the watchlists and how people are added to the watchlists. We also specifically asked about the details of adverse actions taken by staff, to identify any general accuracy issues and especially to identify any apparent privacy impact and biases on Māori, Pasifika, Indian and Asian shoppers (noting that the system itself did not collect information about ethnicity).

107. During each interview, we met with store managers and/or owners, senior security staff, checkout managers, and the store privacy officer (if different). The interviews were also attended by the FSNI Loss Prevention Manager. A full list of stores participating in the trial, and stores visited by Compliance Officers is in Appendix 3.

108. Interview topics included:
    (a) Staff understanding of how FRT works
    (b) How FRT interacts with other store systems
    (c) Number and position of cameras in store
    (d) Signage and information available to customers
    (e) Staff understanding of deletion and retention of FRT information protocols in practice
    (f) Security of staff logons/system access
    (g) Privacy understanding generally
    (h) How staff managed unconscious bias and cultural awareness
    (i) Responding to privacy breaches and incidents.

109. All stores visited operated two or more FRT cameras in-store, including one positioned to capture the entrance foyer.

110. During interviews, all stores reported serious verbal and physical abuse towards staff daily, some of which resulted in time off work for injuries and post-incident stress and anxiety. Some examples of the nature and impact of the serious incidents reported by stores are below:
    (a) Physical attacks on staff causing serious injuries and hospitalisations including punching and kicking and attacks with weapons including baseball bats, screwdrivers, store trolleys and iron bars
    (b) Verbal abuse of staff, including racial abuse and death threats. This abuse is often particularly directed at checkout staff
    (c) High-volume thefts for example, trolleys filled with meat, alcohol or hair products.

111. The stores interviewed reported that since the FRT trial commenced, they had seen a significant decrease in aggressive behaviour and high-volume shoplifting. Some of the interviewed stores had noticed offenders were not returning after being warned that FRT was in operation, or after they were identified for a first time by FRT and told to leave. Where some staff noticed an FRT alert occurring multiple times a week at the start of the trial, by the end of the trial this had reportedly decreased to minimal alerts, with sometimes a week or more going by before another alert went off.

112. Checkout managers and security leads were asked about how they and their staff felt about FRT in their workplace. Both reported that incidents are now less confrontational which is helping staff who have suffered post-traumatic impacts from those experiences. Prior to the trial, when staff challenged offenders, the situation would often quickly escalate and could turn into aggressive and physical altercation. With FRT, staff felt situations could be better managed and de-escalated. Where the offender was considered violent and not approached, staff felt they had more control of the situation by being able to proactively monitor the person instore, promptly contact the Police and provide information for referral to Police in any prosecution that may result. Staff reported a huge improvement in their level of comfort and safety in being at work.

113. Following the store visits, we wrote to each store with commentary on their privacy practices generally and regarding the FRT implementation, including identifying areas of improvement needed.

Changes made during the trial

114. FSNI made some changes to the FRT process during the trial, to address misidentification problems and mistakes in how staff responded. The main changes were:
    (a) removal of low-quality photos from watchlists to reduce the potential for error
    (b) clarifications to the User Manual and further training of staff
    (c) lifting the threshold before staff would act on an alert ("the operational threshold"). The algorithm triggers an alert at a 90% accuracy rate for a match. Originally, staff would be able to intervene at that level ("the operational threshold"). However, after two people were misidentified, the operational threshold was lifted to 92.5%.

115. There were no further incidents of serious misidentification after these changes were implemented.

Variations in trial methodology from initial plan

116. The original trial methodology involved a comparison of up to 25 stores with FRT ("trial stores") with at least 25 comparable stores using the more traditional security measures such as bodycams and security guards, but without FRT installed ("control stores"). Having control stores was intended to provide clear statistical comparisons that would demonstrate whether FRT made a difference to the types or numbers of crime incidents.

117. The report by Scarlatti, the independent evaluator, and responses from FSNI posttrial shows that this methodology was not followed. According to the evaluation report, this appears to have happened as a result of a variety of problems. For example, the sample of stores using the FRT was pre-selected and small, and there were other operational limitations such as unreliable reporting information, unsuitable CCTV and low offence numbers in the non-FRT stores.

118. As a result, the trial became a non-random controlled trial that compared how outcomes changed over time for the stores with and without FRT. It compared the 25 FRT stores with 128 non-FRT stores (including the initially intended 25 control stores).

Data limitations and trial constraints

119. This trial and the resulting data present the following limitations and constraints:
    (a) The total number of unique faces scanned by the FRT systems was not available. Only the total number of images taken by the FRT system

(225,972,004) was available, but since people can be scanned multiple times on a single visit, this number includes duplicates of the same face.

(b) The proxy provided by FSNI – the number of transactions at checkouts – was also not reliable enough to inform a fully objective view of the privacy impact.

(c) The number of images deleted relating to specific individuals was not as clear as would have been ideal. An approximation, through total system deletions of all images or "tracks", was able to be calculated.

(d) Numbers of scanned images, matched images, alerts, interventions and mismatches were not broken down by skin tone, making it hard to evaluate the potential degree of bias for people with darker skin tones.

(e) The trial was constrained by the number of participating stores (25 stores) and the length of the trial (six months). However, these design constraints also ensured the privacy impact of the trial was contained.

(f) Treatment stores were not randomly allocated to the treatment group but based on keenness to participate in the trial. The control group ended up consisting of all non-FRT stores for which data was available. This is not ideal and likely to have introduced some bias to the data.

(g) Reporting behaviour between the two groups (treatment and control) was quite variable. While treatment stores reported consistently throughout the trial, some non-FRT stores became disengaged with the trial over time, resulting in inconsistencies in reporting and comparatively very low incident counts. This affected the overall data quality.

120. In short, it was clear to us that the data need to be treated with a degree of caution. While FSNI and Scarlatti did their best to overcome the data weaknesses, the trial results are not as conclusive as would have been ideal. However, the combination of the evaluation results and our own Inquiry methodology have enabled us to be confident in our conclusions that, overall, the use of FRT was effective at reducing harmful behaviour in the stores in which it was trialled, and that the privacy safeguards worked to protect people and to enable FSNI to comply with the Privacy Act.

121. The next section of this report explains in more detail how we came to this conclusion.

6 Whether the FRT operating model complied with the Privacy Act

IPP1: 'lawful purpose' and 'necessity'

122. The starting point for assessing FRT under the Privacy Act is set out in IPP1:

> *Personal information must not be collected by an agency unless –*
>
> *(a) The information is collected for a **lawful purpose** connected with a function or activity of the agency; and*
> *(b) The collection of the information is **necessary** for that purpose.*

123. In other words, FRT must be a justifiable, effective and proportionate way of dealing with the problem that it is intended to resolve.

Lawful purpose

124. Purpose is a key concept in the Privacy Act's framework. An agency's ability to collect, use and share personal information depends on the purpose for which that information is needed. Our Privacy Act is purpose-driven, rather than consent-driven.

125. This is a major difference between the New Zealand Privacy Act and some overseas privacy laws and is highly relevant in the context of biometrics. For example, under the Australian Privacy Act 1988, an agency cannot collect sensitive information about an individual (including biometric information), unless the individual has consented or an exception applies, such as preventing serious threats to safety.

126. While businesses and other agencies have a degree of autonomy under the NZ Privacy Act in establishing their lawful purpose, it is essential that they make it clear what they are trying to achieve when they collect personal information. The broader or vaguer an agency's purpose for collection is, the harder it will be for that agency to demonstrate that collecting the information is necessary in the circumstances and that they have the appropriate privacy safeguards in place.

127. Being clear about the lawful purpose for using FRT also enables agencies to make sure that they have the appropriate technology, processes and people in place to achieve that end. For example, an FRT system that aims to identify and pre-empt violent crime will be differently designed than a system that aims to verify the identity of staff entering hazardous areas of a workplace. Both the degree and type of privacy risk are different, and therefore the necessary safeguards required to ensure privacy is not unjustifiably impacted will also differ.

Finding: FSNI had a lawful purpose for using FRT

128. FSNI's privacy impact assessment records that supermarkets experienced an increasing number of incidents and breaches of trespass notices by repeat offenders in the period up to March 2024 (the trial commenced on 8 February 2024). Repeat offenders were estimated to be responsible for around one third of retail crime incidents.

129. While all retail crime is problematic for store owners, the specific purpose for introducing FRT was to proactively reduce the numbers of "harmful behaviour" incidents by repeat offenders, either by enabling in-store interventions or by compiling sufficient evidence of offending to support prosecution. "Harmful behaviour" was defined as a range of unlawful and disruptive behaviours including theft, burglary, robbery, assault (physical and verbal) and other aggressive, violent and threatening behaviour. In practice during the trial, stores focussed on more serious harmful behaviours such as assault, verbal abuse and higher value shoplifting.

130. FSNI therefore had a dual purpose. One objective was to reduce the incidence of serious behaviours affecting its staff and customers, such as physical and verbal assault, and aggressive, violent and threatening behaviour. This objective has both a protective element by limiting harm to individuals, and an economic element in providing a safe environment to work and shop.

131. The other objective was to reduce the incidence of high value theft, burglary and robbery by repeat offenders, to reduce the most serious financial losses. There is overlap between the two objectives as the evaluation report notes that most serious events of aggression are also connected to shoplifting.

132. We were satisfied from the start that FSNI had a clearly defined and lawful purpose to use FRT in stores, given the scale and seriousness of retail crime, and the harmful behaviour that FSNI staff and customers experience in stores. That purpose was not over-broad but was clearly focused on areas of most concern.

133. However, it was important to use the trial to test whether the technology was "necessary" to achieve this purpose – i.e. would it work?

Necessity

134. The term "necessity" carries a lot of weight, both in New Zealand privacy law and overseas. It means more than "desirable", "expedient", "reasonable" or "commercially convenient". Instead, the agency needs to be able to show that:

(a) collecting that personal information makes a **clear, demonstrable contribution** to achieving the specified purpose, and the information is relevant and not excessive or arbitrary

(b) collecting the personal information is a targeted, **effective** and accurate way to achieve that purpose: if it does not work, then it was clearly not necessary to collect the information in the first place

(c) there is **no less intrusive option** that the agency could have reasonably and practically used in these circumstances to achieve the same result. If a less intrusive option is available and gets the agency to the same place, it cannot be said to be necessary to adopt the more intrusive option

(d) the **scope** of information collection is also relevant to the degree of intrusion. The more information collected – and the more people affected – the more challenging it might be to show that collecting *all* that information is necessary to fulfil the purpose of collection. The degree of intrusion is also likely to depend on the privacy safeguards that are in place and whether they are effective.

135. In the words of the UK Information Commissioner's Office:

> *To ensure that LFR [Live Facial Recognition] is necessary, [agencies] should be able to demonstrate that LFR allows them to take a particular action and that this requires the collection of biometric data.*
>
> *LFR does not have to be the only possible means of achieving the objective, but [agencies] must consider other alternative measures which are less intrusive and demonstrate that they have discounted them for adequate reasons.*
>
> *[Agencies] should not use LFR simply because it is available, it improves efficiency or saves money or is part of a particular business model or proffered service. While it may be justifiable in some circumstances, if the deployment of LFR is only likely to be slightly more effective than less privacy-intrusive measures (such as non-biometric measures, e.g. alternative types of surveillance) then it may be unnecessary.*

Finding: The trial showed that FRT was an effective way to achieve the purpose

136. Despite some limitations of the trial data, it was clear from both the data and our own inquiries that FRT made a substantial difference in reducing rates of serious recidivist crime in stores.

137. The evaluation results were different for serious harmful behaviours compared to shoplifting. Overall, the evaluation found an estimated reduction in the incidence of serious harmful behaviours of 16% and an estimated reduction in shoplifting of 21%. The majority of shoplifting incidents were stopped by staff, although a significant minority of shoplifting incidents (73 out of 193) were not prevented.

138. The evaluation found a 54% increase in the detection of breaches of trespass, which was a useful indicator that FRT is effective in identifying repeat offenders.

139. The evaluation concluded that FRT avoided around 115 serious harmful incidents over the course of the trial across the 25 trial stores. This equates to approximately 19 incidents per month across all 25 stores.

140. The numbers of incidents avoided were attributed to two categories. One category is the deterrent effect of FRT (around 65 incidents), based on the number of people who complied with trespass notices and did not enter the store. The second category (around 50 incidents) is due to the direct effect of a staff member responding to an alert following a match and a serious harm incident then not taking place.

141. As discussed under 119, the trial data had some limitations and we need to be careful when interpreting these results. Scarlatti, the independent evaluator, cautions to interpret their estimates "*as valid average treatment effects for treated Stores (i.e., 'ATT's), and indicative of potential treatment effects for similarly keen Stores*". This means, these estimates should only be interpreted in a very narrow sense. Making broad generalisations about the overall benefits of the benefits of FRT in retail stores could potentially be misleading.

142. However, putting all the data together with our own observations, we accept that the evaluation demonstrates that introducing FRT made a clear difference in FSNI trial stores, especially as it relates to the most serious incidents involving violence. There was a steep drop in the number of violent incidents during the trial period. Twelve out of 85 incidents included an element of physical violence in the first half of the trial, which decreased to one out of 48 incidents including an element of violence in the second half. That decline was accompanied by an overall decline in the number of FRT alerts over the course of the trial which the evaluator attributed to changes in background offending rates, reduced enrolments of people on watchlists and a deterrent effect from FRT.

143. It is worth noting, though, that increased intervention with customers triggered by the use of FRT can also *create* opportunities for violence. For example, nearly 300 incidents of shoplifting or aggressive or difficult behaviour occurred following approaches by staff over 24 stores, including 13 violent incidents, (none of which caused injuries) and these incidents occurred largely over the first half of the trial. Overall, 17% of alerts actioned involved an incident of some kind.

144. It is certainly important to train store staff in when and how to intervene, and in de-escalation techniques, but it is important not to overlook the crucial role that the Police must still play in responding to retail crime.

Finding: Other less intrusive security measures were insufficient on their own to address the problem but need to continue to be part of the package

145. FSNI already had other reasonable alternatives in place, which are important elements of an in-store security system. However, those alternatives do not seem to be enough on their own to manage the levels of crime. Indeed, the PIA recorded that the number and seriousness of harmful behaviour incidents was increasing. It was logical to test whether it would be useful to add FRT into the toolbox available to stores.

146. However, as discussed earlier, FRT is not a silver bullet. Even if FRT is operating, it is still important to maintain those other safeguards to manage incidents involving first time offenders or others who are not on (or who cannot be justifiably added to) the watchlist.

Finding: The inherent degree of intrusion was very high, but the privacy safeguards employed throughout the system reduced the actual degree of intrusion to an acceptable level

147. The scale of personal information collected through FRT is a major privacy concern, with all customers scanned (usually multiple times) as they enter and move through the store. Of course, capturing such information is inherent in how FRT operates. However, the inherent degree of intrusion (particularly for people not on the watchlist) is extraordinarily high. This means that the privacy safeguards need to be designed in throughout the system to reduce that level of intrusion to an acceptable level.

148. The initial safeguard that made a major difference was the automated and immediate deletion of non-matched images. That design decision meant that the vast majority of

customers could be confident that their information was held only fleetingly, and that it could not be used in any way that would affect them in future.

149. Match accuracy levels are also a significant factor. We set out further details of our expectations about the full package of safeguards required throughout the operating system later in the report.

Accuracy, retention and security issues (IPPs 5-9)

150. IPP8 states that:
An agency that holds personal information must not use or disclose that information without taking any steps that are, in the circumstances, reasonable to ensure that the information is accurate, up to date, complete, relevant, and not misleading.

151. What is "reasonable in the circumstances" will depend on factors such as how the information will be used, and the extent and nature of any risk to individuals. The more serious the effect on people could be, the higher the level of care that agencies need to take to make sure the information is fit for purpose before using it.

152. There are close connections between IPP8 and IPP1. If there are inherent accuracy issues with the design and operation of an FRT system, this will make it harder for an agency to demonstrate that use of the FRT system is necessary, effective and proportionate overall.

153. In the FRT context, accuracy issues most often arise in these situations:
    (a) **Compilation of watchlists** (against which the FRT system matches individuals entering stores), in particular the quality of images and whether a person met the criteria for addition to the watchlist.
    (b) **The accuracy of the algorithm,** including how it performs in relation to specific groups of people.
    (c) The **level at which an alert is triggered**. In the context of FRT, the lower the match level is, the less likely it is that the accuracy level will be high enough to act on.
    (d) The quality of the **human confirmation of the match.** Staff who are asked to validate whether the alert relates to the right person need to be properly trained and supported to do that.
    (e) Any **additional information** that is accessed, used or subsequently recorded about that person, including details of their past behaviour, and a record of what happened if staff intercept them after getting an FRT alert.

154. Accuracy is a key theme of this Inquiry due to the risks of misidentification posed by FRT throughout the operating model. This is therefore a good point at which to consider all the different aspects of the operating model in more detail, to see whether FSNI took reasonable steps during the trial to make sure that the personal information collected through FRT was fit for purpose before it was used.

155. Other IPPs are also relevant to maintaining data quality and fitness for purpose, particularly IPP5 (reasonable security safeguards), rights of access and correction (IPP6 and IPP7) and requirements not to keep personal information longer than is necessary, given the purpose for which it will be used (IPP9).

Elements of the operating model

156. Introducing FRT into an organisation does not begin and end with installing technology. The success or failure of FRT depends on how each step of the whole operating model works. Applying a privacy lens to each of those steps enables a business to identify and eliminate or reduce potential problems that could harm people, undermine the purpose of using FRT, or damage public confidence in the business.

157. This Inquiry therefore reviewed the end-to-end operation of the FRT operating model during the 6-month trial period, including considering the results of the independent trial evaluation. Broadly speaking, an FRT operating model falls into the following areas:

    (a) **Key considerations before deploying FRT**, including clearly identifying what the purpose of introducing FRT is, considering whether there are other options for achieving the same end, selecting the appropriate technology, and conducting risk assessments (this is dealt with earlier in this report under “lawful purpose”)

    (b) **Setting up and maintaining the watchlist,** including deciding who needs to be enrolled to meet the purpose, avoiding bias, whether appropriate quality images are available, and how long a person remains on a watchlist

    (c) **Installing and maintaining the wider FRT operating system**, including monitoring for biases and inaccuracies

    (d) **Information for customers about the operation of FRT**, and how they can access or correct information that the store holds about them.

    (e) **Alerts**, for instance setting an appropriate level at which an alert is triggered, and the human verification process (including training for staff about how to read the results)

(f) **Decisions about whether to intervene**, including ensuring appropriate staff training, how to respond to incidents involving people aged under 18 years of age and vulnerable people, appropriate involvement of Police, and avoiding bias.

(g) **Managing complaints and access requests under IPP6**

(h) **Review and monitoring**, to ensure the system is still functioning as intended and is achieving the desired results.

Watchlist criteria and practices

158. A diagram of the information flows and operational settings of the watchlist processes is in Appendix 4. The total number of watchlist enrolments peaked at almost 1,800 at the beginning of the trial. At the end of the trial there were 1,504 enrolments across the 25 trial stores.

159. The store watchlists are a crucial component of the FRT system as they set the parameters for who is targeted by the FRT process, and the reasons for being targeted. The quality and accuracy of the watchlist is key to the success of the FRT. Most of the IPPs are relevant to how watchlists are created and managed.

160. Our key expectations when deciding whether to add someone to a watchlist are:

(a) Adding that person must be clearly relevant to fulfilling the purpose that FRT is intended to achieve.

(b) The information on the watchlist needs to have been collected in a way that was fair and reasonable in the circumstances.

(c) Objective and consistent criteria for enrolment are needed to mitigate the risk of subjective decision making that could perpetuate unfairness, bias or discrimination.

(d) Enrolment decisions must be managed by a small and well-trained group of people.

(e) Children and young people or other vulnerable people should not be added to a watchlist.

(f) People should only be added to a watchlist based on objectively verifiable facts (such as a conviction, clear evidence of relevant behaviour, or a trespass notice). Taking adverse action against someone on the basis of information that is opinion based carries more risk – both for the person concerned and for the agency – than information that has been verified.

FSNI's watchlist criteria

161. An offender or accomplice is added to the store watchlist by an authorised staff member where that offender or accomplice meets one of the following criteria:
    (a) They have engaged in, or contributed to, serious harmful behaviour (physical assault or verbal abuse, behaving in an aggressive or threatening manner, damaging store property)
    (b) They have engaged in, or contributed to, high value shoplifting
    (c) They are a repeat offender, unless the theft is of a very low value
    (d) They have been trespassed from the store.

162. Individuals suspected or identified as "vulnerable", or who may be under 18 years old are not entered into the watchlist. If staff are unsure of the person's vulnerability status or age, they err on the side of caution and the person is not enrolled.

163. Decisions to add someone to a watchlist are one major point at which bias or discrimination can be introduced or existing bias amplified. This is something that needs to be checked regularly. However, our compliance team specifically reviewed the watchlists during store visits and were satisfied that there was no apparent evidence of discrimination at that point during the trial.

Finding: **FSNI's watchlist criteria are appropriate, targeted, and sufficiently protective of vulnerable people. There was no apparent evidence of bias or discrimination**

Establishing the watchlist

164. Each FRT store created and maintained its own watchlist and information contained within that is not shared with any other party, including any other FSNI store.

165. Before the trial started, two authorised staff from each FRT store reviewed that store's existing watchlist in the store's incident reporting system and then uploaded images and notes about individuals who met the FRT-specific criteria into the FRT watchlist. After that, no more images were transferred from the incident reporting system to the FRT system.

166. The information uploaded to the FRT system included the incident reporting reference number, but there is no direct link between the incident reporting system and the FRT system. Manual transfer helps to ensure that the information is checked and a judgement made by a trained staff member before information is added to the FRT watchlist.

167. Sourcing the initial information from the store's incident reporting system meant that the accuracy of the original dataset of information on the FRT watchlist relied on the accuracy of the information stored on that system. The normal criteria for adding information to the platform are broader than the criteria for adding someone to an FRT watchlist. There is a particular risk that the existing incident reporting information is sometimes more subjective in nature, which could make it unsuitable for use on an FRT watchlist. Again, ensuring a trained staff member makes the decision is a useful protection against adding information that goes beyond the permitted purpose of FRT.

168. We note that the evaluation of the non-trial stores was constrained by the quality of the data held in their incident reporting system. We **recommend** that if the store uses that system, store staff improve the quality of the data both to support normal use of information in that system and so that any future watchlists are more accurate if the incident reporting data is used as the base in future. This recommendation was also noted by Scarlatti in its evaluation report.

169. There were two serious misidentification incidents early in the trial. On review, one contributing factor was low quality images on the watchlist, which affected the accuracy of the match process. The required image quality was then improved and the watchlists were cleaned up. Other changes were also made and no similar incidents were reported after that.

170. While it was positive that the errors were addressed, the incidents were acknowledged to be harmful for the people who were misidentified. Since image quality is well known to be one core factor in how well an FRT system performs, we consider that FSNI should have checked for low quality images before the trial went live.

Finding: images that were too low quality were initially included on the watchlist. This was one contributing factor to the harmful misidentification incidents that occurred. FSNI subsequently reviewed and improved image quality

Maintaining the watchlist

171. The watchlist does not remain static. Following an incident of harmful behaviour, an authorised staff member will update it with images of a person of interest or update existing information about that person.

172. During the trial, this process required two authorised personnel to confirm that the individual is a person of interest (either an offender or an accomplice) and that the enrolment information is accurate. They also needed to reasonably believe, based on

supporting evidence, that the relevant enrolment criteria apply. A subjective belief was not enough. Staff were trained in how to record incidents and behaviours.

173. Our compliance team observed that store staff are currently very cautious about enrolling people on watchlists and cautious about getting it wrong. As with any new practice, this cautiousness may fade over time, so it is crucial that FSNI retain a long-term focus on auditing watchlist enrolments and false matches.

Finding: staff were adequately trained to add or update information on the watchlist

Watchlist information retention and review

174. FSNI set watchlist retention rules so that enrolment expires after a maximum of 2 years for principal offenders (to equate with the period of a trespass notice), or a shorter period selected by the store based on the incident. The period is three months for accomplices: information about whether someone is an accomplice can be inherently more subjective, and they are not subject to such lengthy exclusion periods.
175. Store privacy policies set out the process for an individual to request removal of their enrolment on the watchlist and review by the store. People can ask for their information to be removed if they consider that they have been incorrectly identified or did not meet the enrolment criteria or there are other circumstances to consider. If an individual is not satisfied that their correction request has been adequately actioned by the store, they have a right to complain to OPC under IPP7 of the Privacy Act.

Finding: FSNI has set clear limits for how long people can remain enrolled on a watchlist.

176. Those limits are appropriate for people who have been trespassed as they accord with the length of the trespass notice. However, a retention period of 2 years could be excessive if the person had engaged in less serious behaviour. We **recommend** that retention periods should be less than 2 years if the behaviour is at the less serious end of the spectrum.

177. We also **recommend** that it is important for FSNI to schedule regular reviews of image quality, enrolment decisions and other key discretion points, so quality of decisions do not deteriorate over time.

Controls and security relating to the watchlist

178. Our interviews with staff and on-site visits confirmed that key security settings are in place (both to protect the integrity of the watchlist and other operational elements of the system):
    (a) Only authorised store personnel have access to the FRT system both in terms of logins to the IT system and access to the security room.
    (b) All access is logged and regularly reviewed by the FSNI Loss Prevention Manager.
    (c) The FRT system is not connected to any other in-store system.
    (d) FRT alerts can only be received by authorised devices, and those devices only work on the in-store network, so are ineffective outside the store site.
    (e) The FRT screen was positioned so that it was not accessible from the doorway of the security room – that is, it was tilted away from the entrance, and the security room doors were closed when the system was being used.

179. Scarlatti's report also noted that these system security settings were maintained during the trial and no security breaches were reported.

180. We note that it is not sustainable to continue to place sole responsibility on the Loss Prevention Manager for reviewing and managing the watchlist. It would be advisable to train one or two back-up personnel. Having only one authorised person in this role means that any changes required for user accuracy, role-based access control and the enforcement of internal access policies cannot occur if that person is not available. Such gaps can lead to security issues. Similarly, having more than one store champion in each store would ensure that a person well trained in FRT protocols is always on site.

Finding: Information security controls during the trial were robust, which is appropriate given the sensitivity of the information involved

Findings relating to watchlists, and recommendations for further improvements

181. We are satisfied that the current watchlist settings and management processes are reasonably well designed. While accuracy-related issues led to some misidentification incidents, these resulted in system changes and improvements being made during the trial.

182. However, from the trial we have identified several measures that would strengthen watchlist processes in future. Also, existing settings can easily change when a practice becomes business as usual, and personnel change. Taking the following

additional steps would let FSNI be more confident that store watchlists will operate consistently and that privacy safeguards will be effectively implemented over time:

(a) **Clarify criteria for incidents involving attempts to steal and the level of supporting evidence required for the individual to be enrolled**. Store protocols define that a person may be enrolled if two authorised personnel believe, on reasonable grounds based on supporting evidence…" that harmful behaviour has occurred". There is not defined specific criteria for what supporting evidence is required. It should be clear that suspicion alone is not sufficient to meet watchlist criteria.

(b) **Clarifying the criteria for listing someone as an accomplice.** Currently an accomplice is defined as someone who "actively assists but were not the instigator of the incident". For consistency across watchlists and stores, ensuring staff have the same protocols and understanding of this definition is critical.

(c) **Criteria for theft should specifically relate to high value shoplifting** as described in the evaluation report and exclude very low value shoplifting as a ground for enrolment. This would formally mirror current practice. In interviews with store staff, it was observed that security staff applied a level of discretion as to whether enrolment for a 'minor' offence warranted enrolment in the watchlist.

(d) **Clarify that a trespass notice is a basis for enrolment only if the reason for issuing the trespass notice is consistent with the enrolment criteria** (that is, the person has engaged in harmful behaviour).

(e) **Update and strengthen the references to watchlist criteria in the PIA** – there is currently a discrepancy between the operational practice and the PIA in some respects.

(f) **Continue to review watchlists regularly** – conduct human checks to complement the automated deletion rules, to ensure that watchlists do not contain out of date or biased information. That includes capturing data on misidentifications, including skin tone, to understand and remedy the underlying cause.

Technology selection

183. In choosing a vendor for the trial, FSNI relied on assessments from the National Institute of Standards and Technology (NIST), an agency of the US Department of Commerce that evaluated the accuracy and demographic bias of the Imagus FRT System. NIST tested the FRT System against its global database of facial images and considered that the Imagus FRT System was the second best performing FRT system in the world for processing images "in the wild" at that time. Images "in the

wild" are those where the subject does not pose for the image (for example, images that are taken from CCTV footage, or images taken where people are on the move). This quality control was therefore essential for the context of FRT in supermarkets, where people do not pose to have their face scanned and where (unlike in gaming venues) security staff would not ask people to remove hats, masks and so on, and pose for a camera.

Finding: FSNI chose an FRT system that had been professionally tested as suitable for operating "in the wild"

184. We understand that the Australian-based system developer has continued to diversify its database to include images of groups including Māori and Pacific peoples. However, as mentioned earlier, we remain concerned that there is no FRT system currently available that is trained for the unique features of the New Zealand population. It is therefore hard to be confident about the level of accuracy that the system can achieve in New Zealand conditions.

185. The evaluation was not able to provide any conclusions about this risk with any certainty due to the lack of data on how the system performed for people with different skin tones, and a lack of relevant research. The evaluation noted that the potential for bias was mitigated in the trial by the two-person verification process. This was considered to provide a check against any human bias (at the relevant points of discretion) as well as any bias of the FRT system.

186. However, there are potentially amplified risks if machine bias is present, and human correction will not necessarily address them all. The likelihood of privacy interferences is disproportionately increased where FRT has an inbuilt bias, as explained in a paper by the Alan Turing Institute on bias in facial recognition technologies.

187. We are satisfied that the ability of the technical system and processes to mitigate against bias were probably as good as could be achieved with the current lack of a New Zealand training data set. However, as more entities consider using FRT, we consider it is becoming pressing to develop a local training dataset, so that the public in general, and affected minority groups in particular, can be more confident that bias has been reduced or eliminated in the system.

188. The images captured as part of the trial were expressly *not* allowed to be used for training purposes. That was – and remains – an important privacy safeguard, as it would exceed the purpose for which the systems were being used, would result in

excessive retention of images that were captured non-consensually, and may have raised significant data sovereignty and jurisdictional issues.

Match accuracy and alerts

189. Whatever the quality of the FRT system, accuracy is never entirely assured. FSNI therefore created further mitigations. All the protocols for verification of alerts and use of devices are set out in the FRT User Manual and Zebra Device Manual (that is, the handheld device used by store staff), as well as other protocols set by FSNI:
    (a) First, FSNI configured the system to ensure that a match must achieve an accuracy rating of 90% or more before an alert is triggered.
    (b) This alert is then checked again through the in-store process where the alert is reviewed and the match verified, providing an assurance to determine the accuracy of the match before any action is taken by staff.
    (c) An FRT alert is not actioned unless an alert is generated on more than one camera. This increases the validity of any possible matches and is an important safeguard. The multiple cameras also reduce the possibility of people of interest being missed due to a face being obscured, for example by a hood or mask.
    (d) Originally, staff were able to act on an alert provided that the accuracy was 90% or above. However, following the misidentification incidents, this was adjusted. The FRT User Manual now sets an operational threshold of 92.5% accuracy confirmed by two or more FRT camera alerts before staff can consider whether to act on it. Staff were instructed only to observe the person if the match score was between 90% and 92.5%.
    (e) Staff were trained to look for the appropriate accuracy rating before determining whether to act. Staff were also made aware that the system could return errors, so their human judgement was important.
    (f) Some stores set a higher level at which staff would act after an alert (the highest was 98%). None were under the FRT User Manual stipulation of a minimum 92.5% match before action would be taken. The variance was a reflection that each store could choose to implement the technology at a level that it considered to be appropriate for its own community.

190. Setting the minimum accuracy match is a finely balanced decision. The higher the match accuracy is, the less likely someone is to be inconvenienced or embarrassed by being incorrectly approached as a person of interest. However, setting that match accuracy too high can lead to automation bias – that is, where staff assume that they can simply rely on the computer without using their own judgement. It is essential for all processes – and staff – to recognise that a match score is just one tool to help to

decide what to do next and how best to handle a situation. Acknowledging that there is a margin of error makes it less likely that staff will fail to listen when someone explains that there has been a mistake.

191. However, we consider that the trial has provided evidence that it would be preferable to set a higher match score to avoid misidentifications, and that setting that higher score will not compromise the ability of FSNI to fulfil the purpose for which it is using FRT. In their evaluation, Scarlatti also suggested that FSNI should take extra care if the match score was below 94%. This could be the new operational threshold.

Finding: **The algorithm should be adjusted to trigger alerts at a minimum of 92.5% to reduce risks of misidentification. FSNI or individual stores can then choose to set a higher operational threshold for action**

Devices on which alerts are viewed and checked

192. When a potential match for a person of interest is detected, an alert is sent to a handheld FRT 'Zebra' device (about the size of a mobile phone), which is set to only operate on the in-store network. The authorised staff member using this device verifies the alert with another authorised staff member to confirm it.

193. The device carried by staff in store can receive and edit alerts and add notes to the FRT watchlist. The device cannot upload individuals to the watchlist as this is only done from a secure computer in the store's security room.

194. Most stores currently view and compare images from the FRT watchlist and alerts on devices similar in size to a smartphone. Both images are cropped to show an extreme closeup of the face only, and do not include hair or neck or other features outside of the face. Staff reported that this restriction makes it more difficult to locate and correctly identify the individual in a timely way so that interventions can occur before harm occurs. Focusing on the face alone also makes misidentifications more likely.

195. The FRT device limits visibility and may omit key identifying features. In contrast, CCTV footage on larger computer screens provides a wider field of view, capturing details such as the person's clothing and headwear. This significantly improves the ability to accurately identify individuals.

196. We **recommend** that FSNI consider extending the size of the image available to staff to provide enough details of their clothing that they can be promptly and accurately located. This recommendation was also made by Scarlatti. Where the extended

image also captures images of people who are not on the watchlist, their images should be redacted.

197. Scarlatti also recommended that the device provide a real time notification to authorised users where a misidentification has occurred, or a process step has not been followed (for example that there was a ping on only one camera). This would provide a prompt escalation to the store champion and FSNI Loss Prevention Manager that an error has occurred. We agree that it is important for errors and misidentification incidents to be promptly identified and assessed so that the root cause of the error is identified and remedied.

Intervention decisions

198. Of course, while it is essential to have a 'human in the loop', humans are not infallible either. It is entirely possible – even likely – that the staff checking the alerts will make a mistake from time to time and misidentify a person. It is also possible that staff will act in a biased or discriminatory manner, though we did not witness such behaviour during our store visits. Those types of mistakes are not exclusive to situations where FRT is used.

199. Trial design did not allow the evaluator to measure rates of human error leading to false approaches in the control stores, or in the trial stores prior to using FRT. As a result, there is no baseline of what is normal in the absence of FRT. The evaluator's best judgement is that the raw increase in the number of approaches means that it is most likely that FRT *increases* the likelihood that people are falsely approached, even if that likelihood is low. Overall, there is limited information available to confirm instances of misidentification. Due to the information gaps, this is a potential area of concern under IPP 8 that should be kept under review.

Finding: All instances of misidentification should be clearly logged along with records of how the mistake was rectified. Misidentification incidents and responses should be analysed to ascertain root causes, including skin tone, and any reasonable adjustments made to reduce risks of recurrence.

Retention of non-confirmed matches

200. As mentioned earlier, all images of people which did not trigger a match at the requisite level (currently 90% within the system) are deleted within 59 seconds.

201. All other alert data is deleted from the FRT system at midnight on the same day. Data held in the FRT watchlist is held for the period of the relevant individual's trespass notice.

202. This prompt deletion is a privacy protective measure, and we support it. Of course, the trade-off is that there will be times when a record is deleted despite potentially being a valid alert, and Scarlatti's evaluation recommended that the information be retained for up to five days. However, if the person did not commit an offence while in store, it does not matter that someone *might* have been able to ask them to leave. On that basis, we disagree with Scarlatti's recommendation. Instead, we consider that the midnight deletion is an appropriate retention setting.

Creating records of what occurred

203. Following action taken on an alert, staff add notes into a free text field on the FRT system as to what happened, for example, a description of the behaviour of the individual in store, whether the individual was approached, and whether the individual was violent or non-violent. These notes can be viewed in the event of future alerts.

204. We observed that the free text field auto-populates with the notes that have been previously added by the staff. That feature introduces a privacy risk where information about one person could be selected for entry into the notes of another person giving the possibility for inaccuracy of personal information. We **recommend** removing the auto-fill function on the FRT notes field to reduce the likelihood of the previous note entry being added to the wrong file. The evaluator also identified this as a privacy risk and included a recommendation in its report that these free text fields are replaced with drop down or more targeted text fields. Scarlatti also recommended implementing a process that prompts authorised users to add missing information, with the intention that this would improve the accuracy of the notes.

Training

205. Store staff assigned as authorised FRT users received privacy training on the use of the FRT as provided by FSNI. The store champion and security staff were trained in person, and training was also provided online more broadly and supported by documented guidance. Training was focused on store security staff who are trained in person by the FSNI Loss Prevention Specialist. User guides have been written in a way that is easy to follow, with clear flow charts.

206. We consider that the training provided was successful and that staff were appropriately privacy-aware. However, we **recommend** that it is important to refresh the training from time to time, preferably based on real-life examples in stores. Any incidents of misidentification, or insights learned from complaints, can also be useful subjects for updated or refresher training.

207. Our store visits and interviews confirmed that these processes were understood and followed, and staff generally demonstrated strong awareness of privacy. However, three stores that did not fully meet our expectations. In two stores the Privacy Officers confused the privacy procedures for privacy breaches with information access requests, and one store did not have a privacy officer until the day of interview. Again, these findings support the need for refresher training as well as continued privacy support from FSNI headquarters.

Additional privacy safeguards developed by stores

208. We noted that some individual stores implemented additional privacy protections above those set out by FSNI including:
    (a) In one store, the security room was itself monitored by CCTV.
    (b) Staff log out of the FRT system immediately when they finish using it, limiting the risk for unauthorised access or use.
    (c) The FRT screen quickly enters a 'timeout' feature to go dark to limit unauthorised viewing of images.
    (d) In one store, information is transferred from the FRT system to the incident reporting system via a USB stick that is not removed from the room. Information on the USB stick is manually deleted once uploaded.
    (e) We saw examples of double-sided signage that covered multiple points of entrance into store, that is carparks and main entrances.
    (f) Some stores have designated areas for verifying a match alert out of sight of customers.
    (g) Staff apply discretion and using de-escalation tools with customers with a history of less serious offending rather than FRT watchlist enrolment.
    (h) Some stores take a cautious approach to match alerts with a lower than 95% accuracy.
    (i) Some of these measures come with the caveat that each alert or incident is different, and in many cases, staff will apply other contextual knowledge to a situation such as personally knowing a person of interest.

209. The protocols and manuals represent the minimum acceptable standard, but many of these additional protections are closer to best practice. Stores can apply higher standards if that is what is appropriate for them (including declining to use FRT at all).

When the identification process went wrong

Incident at New World Westend, Rotorua

210. The most publicised misidentification incident occurred in April 2024, which was early in the FRT trial. Media reported that a "Māori mum was misidentified as a thief" at New World Westend in Rotorua which participated in the FRT trial. Both FSNI and the store undertook an investigation into the circumstances in that case. The person did not complain to OPC and so we have not conducted our own investigation into the incident.

211. As part of the store visits, we met with store staff and reviewed the instore protocols in place at the time and reviewed the investigation report provided by FSNI. The FSNI investigation noted that this incident occurred early in the trial and the following factors contributed to the misidentification incident:
    (a) Key staff trained in the use of FRT and operational protocols were not on site at the time of the incident. Staff responding to the alerts from the FRT system were not sufficiently trained to assess the accuracy of the match identified by the system.
    (b) The image that was loaded on the FRT watchlist was of poor quality.
    (c) The FRT alert was generated by only one of the four FRT cameras in store, which returned a match with 90.54% accuracy.

212. In response, FSNI implemented further training in that store focusing on:
    (a) The seriousness of misidentifying people.
    (b) Additional checks are required when the match is so close to the alert level.
    (c) The requirement for the FRT alert to be generated on more than one FRT camera to reduce the likelihood of misidentification.
    (d) Any low-quality images (i.e. not of passport standard) have been identified and deleted from FRT watchlists.
    (e) The camera frame rate was upgraded to improve the clarity of images collected.

213. The insights from the incident were shared with all trial stores and the key messages above reiterated. FSNI updated the FRT user manual to reflect that staff should not intervene unless a match is at least 92.5% and should monitor only if the match is 90%-92.5%.

214. We note that this incident occurred in the early phase of the trial and that changes were implemented in FSNI procedures to reduce the likelihood of it happening again.

Other incidents

215. There were up to an estimated 13 instances where the alert verification process failed and individuals were approached without at least two authorised staff members manually verifying the match, or where the staff discovered after the fact that they had the wrong person.

216. Of these 13, nine instances involved people being misidentified. The two most serious incidents were:
    (a) The misidentification incident at New World Westend in Rotorua
    (b) Another incident where a match was identified to be incorrect only after the person had left the store after being intercepted. That person received an apology from the store and was offered a voucher. In this instance the person accepted the apology and "expressed gratitude that the team were learning from their mistake."

217. Apologies were given by the stores in the other seven cases.

218. A further four incidents involved individuals being approached in store without a correct match being confirmed. This does not necessarily mean that they were misidentified: these incidents appear to have simply been failures to follow the correct protocol. Such failures create risks for the store, but, as it happens, it is unclear that harm resulted.

219. The misidentifications occurred due to human error and/or uncertainty during the verification process. These errors occurred in the earlier stages to midway through the trial period with the most recent instance in late June/early July 2024.

220. Overall, the number of recorded misidentifications is low with nine cases out of 1735 total FRT alerts. The lack of such cases at the later end of the trial may reflect the improvements made to operational protocols, photo image quality and staff training as the trial progressed.

221. However, it is important to note that the number of misidentifications could be higher. Store notes for a further 70 cases were unclear about whether the person was a match and whether they were approached in store. A further 31 cases were clear that the person was approached but notes did not record the match. Notes for three further incidents were unclear about the status of the match and whether the person was approached.

222. We **recommend** that FSNI needs to emphasise the importance of keeping good records on actions that are triggered by FRT use, so it can be improved and its utility monitored over time.

223. The operation of the trial in stores otherwise followed the requirements set out in the FRT User Manual. FSNI's approach to the trial was to take a continuous improvement approach, so that incidents and errors that occurred were not repeated. While this does not insulate them from liability if errors do occur, it demonstrates a culture of good privacy practice.

What happens if there is a misidentification

224. FSNI intends to have a process in place for staff to follow if a misidentification takes place. This process involves apologising to the individual, notifying the store champion, updating the alert event notes describing the incident, informing the Loss Prevention Specialist, and providing a full written account to FSNI and the store manager.

225. While we agree that such processes are appropriate, we are mindful of the potentially significant impact on individuals that will be the subject of misidentification. There may be situations where an apology and a voucher is insufficient to address the harm that the person has suffered. It is also important to let people know that they have the right to complain to OPC.

226. As mentioned earlier, while the percentage of misidentifications may be small, rolling FRT out at scale would mean that large numbers of people would be misidentified. FSNI research that supported its public survey estimated that, out of millions of shoppers, around 900 shoppers would be stopped every year after being incorrectly flagged by FRT and asked to show ID or explain who they are before being allowed to continue shopping. While it is not evident in the evaluation report how these numbers were calculated, it is a useful scenario to illustrate the potential impact on people who are not in fact on watchlists.

Recommendations relating to misidentification

227. This is an area that FSNI will need to continue to monitor, given the potential for misidentification incidents to occur. We make the following recommendations for continuing improvement of policies and processes, to ensure that the risk of misidentification is minimised:

(a) the algorithmic accuracy setting should be set at 92.5% (up from 90%);

(b) processes should be included in the user manual for directly verifying the identity of customers who are approached that facilitates the confirmation of

identity with the customer and allows for identification mistakes to be corrected where possible before the intervention is escalated to the customer being asked to leave the store;

(c) there should be ongoing staff training on processes to avoid misidentification incidents. Given the risks of misidentification are potentially higher for Māori and Pasifika customers, we recommend that staff training includes awareness of unconscious bias;

(d) there should be routine assurance activities, including a focus on misidentification incidents and false matches to inform the periodic review of accuracy settings and FRT processes and procedures.

Transparency for customers about the operation of FRT (IPP3 and IPP4)

228. IPP3 requires agencies to notify people of a range of matters at the time information about them is collected. This includes the fact that the information is being collected, the purpose for which it is being collected, whether it will be passed to another agency, whether the person has a choice in the matter, and their rights of access and correction.

229. IPP4 requires agencies to collect personal information only in a way that is lawful, fair and not unreasonably intrusive in the circumstances.

230. Many aspects of lawfulness, fairness and level of intrusion have already been discussed in this report. However, fairness also often hinges on the level of transparency about how information is being collected. Covert use of FRT is inherently unfair and intrusive. In a supermarket context, it was therefore essential for FSNI to be clear about the fact that FRT was operating and why.

231. FSNI used a range of channels to provide information to customers about the use of FRT instore so they were informed that the information was being collected, the purpose of that collection and how it would be used, and how customers could request access to or a correction of that information:

(a) there had to be A1 or A0 signs that were easily visible at all entry points

(b) there were smaller signs within stores.

232. Our compliance team confirmed that these signs were present during our on-site visits and through other evidence provided by FSNI.

233. FSNI also communicated about the trial through media and on FSNI's website. Customers had access to the store privacy policy and information about the trial at instore customer service desks. In at least one store we visited, key staff members and the customer service centre held business-card sized information cards which directed individuals to the FSNI privacy team's email address and some information about customer privacy rights and the FRT trial.

234. FSNI staff interviewed had a good awareness and understanding about the trial so they could answer questions from customers about the technology and how to access their personal information or make a privacy complaint. General store staff (ie those not directly involved in the FRT implementation) had limited information on the trial to answer customer queries. The policy in that situation was for those staff to direct customers to the customer service desk.

235. Of course, these communication mechanisms are not failsafe. Despite the signs and other information, and the context of a high-profile trial, the data provided in the evaluation report showed that only 67% of surveyed customers were aware of FRT being used in their store. Both OPC and Scarlatti consider this indicates a need for stores to ensure a higher percentage of customers are fully informed about the use of FRT and the consequences of the collection *before* they enter the store. Additionally, it would be difficult for a person with low vision or blindness to read signage. To ensure this section of the customer base are fully informed, audio messages in the store foyer could be played.

236. Stores and FSNI should consider further communication activities to increase the percentage of customers that are aware of the use of FRT technology in store, particularly for low vision or blind people, and should continue to monitor levels of customer awareness that FRT is operating.

Finding: While the communications to customers broadly complied with IPP3, improvements are possible.

Complaints

237. FSNI received a total of 155 complaints and enquiries through different channels. Our review of this anonymised feedback gives the following insights:
    (a) Most individuals who complained felt their privacy was going to be breached by FRT and that collection of their biometric information was invasive. These individuals felt they were under surveillance, and this made them feel

uncomfortable. They did not accept the use of the technology and were not going to visit stores that have FRT in operation.

(b) Some individuals indicated they would like to have an option to opt out of being monitored – they felt being monitored while shopping for essential items as violations of their rights to privacy and right to free movement.

(c) Some were interested in transparency and clearly knowing which stores were part of the trial so they could make an informed decision where to shop going forward and expressing a choice of not shopping at stores that had FRT in operation.

(d) A minority sympathised with FSNI in trying to find a solution to retail crime, but felt the FRT was not warranted. These individuals were concerned about the untested nature of the technology, and that this could cause issues for people of other ethnic groups such as biases and mistakes.

238. OPC received 123 comments from the public during the trial. 104 were not supportive of the trial, ten were neutral, and nine were supportive. These comments had similar themes to the feedback given to FSNI. For those that were supportive or neutral, they expressed a view of FRT being a good safety tool.

239. Although we did not receive any individual complaints during the trial, any person who was misidentified, either during or following the trial is able to complain to OPC (if they are not able to resolve the complaint directly with FSNI).

240. Concerns about the use of this technology are unsurprising. The indications that they feel uncomfortable with it are also especially important. These perceptions are a highly relevant factor for businesses that wish to retain their customers' trust. FRT is not a simple solution that everyone will readily accept.

241. However, it is worth noting that FSNI's own public opinion surveys suggested that a majority of people were broadly comfortable with the use of FRT to address retail crime as it operated in the FSNI trial, including the privacy protections that were used in trial. On the other hand, a majority of people preferred more security staffing to FRT in principle, as long as it did not affect grocery prices.

7 What other retail businesses can learn from this Inquiry

242. As well as supporting FSNI's use of FRT, the FSNI trial will be useful for other retail businesses considering using FRT to combat retail crime (or for other purposes). As will be clear from this report, FRT is inherently intrusive, and any use of it requires strong justifications and careful system design to ensure all appropriate privacy safeguards are built in.

243. Our key expectations for businesses include the following:

(a) **Careful thought before deploying FRT.**

An FRT business case should consider:

- What is the specific problem you are trying to solve (your purpose) and how serious is the problem?
- What options are available besides FRT as a response and how do they compare from an effectiveness perspective?
- Will FRT make enough of a difference to justify the cost of implementing an end-to-end FRT system?
- If FRT looks like the most effective option, is it a proportionate way to address your problem?
- What privacy risks will you need to manage, and can these risks be managed or mitigated?
- If you operate across multiple sites, are there differences between them that might affect whether FRT is appropriate in each?
- How will you protect the particular interests of children and young people, or other vulnerable people?
- What system will you use and has it been risk tested?
- Do you need to trial the FRT first to check whether it's worth using, or to test your processes?
- Can the system automatically delete all people who don't match your watchlist?

(b) **Setting up and maintaining the watchlist**

Key questions a business should ask are:

- How will you define your target group and set clear, justified and consistent parameters for your watchlist?
- What is your source of images for your watchlist? How accurate is that information?
- Is the image quality fit for purpose?
- How will you check whether your watchlist is biased?
- Who can make decisions about whether to add someone to a watchlist and are they properly trained?
- How will you ensure you exclude children and young people, or other vulnerable people from your watchlist?

(c) **Installing and maintaining the wider FRT operating system**

Key questions a business should ask:

- Have you done due diligence on the software provider's privacy and security policies to ensure the required standards can be met, and if the provider will use the data for training purposes?
- How many cameras will you have and where will they be placed to achieve your purpose?
- What accuracy setting to choose to strike the right balance between false negatives and false positives?
- What data will you need to collect to monitor that your system is operating as intended?

(d) **Information for customers about the operation of FRT**, and how they can access or correct information that the store holds about them.

Key questions a business should ask:

- What type of signage will be effective communication for your premises (before entering, and in-store)?
- What information can you put on the signs to make it clear FRT is operating, without overwhelming people with detail?
- How will you provide people with more detail if they need it?
- How will you deal with customer queries, concerns or complaints? Will your staff know how to respond?

(e) **Alerts**

Key questions a business should ask:

- How will you ensure there is human decision making to check the alerts?
- How many people need to check that an alert is accurate?
- Will you need more staff to run your FRT alert system?
- What information and training do your staff need to carry out their assigned roles? How will you make sure they know the system is not failsafe and they need to exercise judgement?

(f) **Decisions about whether and how to intervene**, including ensuring appropriate staff training, appropriate involvement of Police, and avoiding bias.

Key questions a business should ask:

- What training do your staff need?
- What policies and processes will you need to support staff and customers?
- Do you have clear objective intervention criteria that avoids the risk of bias?
- What process will you use to verify customer identity in case of an alert error?
- How will you respond to incidents involving people under 18, or other vulnerable people?
- Are your staff clear about when they can intervene and when to call the Police?

(g) **Security of your FRT systems and the sensitive information on them**

Key questions a business should ask:

- Who needs to have access to the FRT system? What do they need access to?
- How will you prevent other people from accessing or modifying the information on the system? Is it stored in a secure room or other restricted area?
- How will you prevent people from extracting it and sending it to someone else, or using it for wrongful purposes?
- What protections are in place to prevent the system being hacked?
- What network do any hand-held devices work on, and will they only work in-store?

(h) **Managing complaints and access or correction requests**

Key questions a business should ask:

- How will you deal with customer requests for information about them?

- How will you deal with customer requests to be removed from a watchlist?
- How will you deal with customer complaints under the Privacy Act?

(i) **Review and monitoring**, to ensure the system is still functioning as intended and is achieving the desired results.

Key questions a business should ask:

- How will you keep records about FRT so that you can review and change your processes?
- What audits and reviews will you need to run to look for unexpected outcomes?
- How will you provide assurance to the regulator and your customers that key privacy safeguards, such as immediate deletion of data, remain in place on an ongoing basis?

244. Being able to answer these types of questions (and others raised by this report) will mean that businesses wanting to use FRT in a retail context are better placed to meet their responsibilities under the Privacy Act and their customers' expectations.

8 For the wider system

245. The FSNI trial provides an opportunity for government and the wider system to reflect on the opportunities for research and other support that would assist with safe use of FRT in New Zealand.

A New Zealand training data set?

246. The lack of a New Zealand training/testing data set means that the risk of system bias with respect to the New Zealand population cannot be fully quantified. While steps have been taken in the FSNI trial to reduce the risk of bias and inaccuracies, it would be highly desirable to explore options for ensuring that New Zealand businesses have access to FRT software products that are evidently suitable for our population. OPC supports opportunities for consent-based research that help ensure that FRT is fit for New Zealand conditions.

A centralised system?

247. A suggestion posed by the evaluation is whether the retail sector should develop a more centralised system including a centralised offender dataset or watchlist. The suggestion is that this may potentially improve the effectiveness of retail use of FRT (where repeat offenders from other locations are not included on a store's watchlist). There are also suggestions that a centralised system could mitigate security risks such as data breaches, based on the assumption that it would be easier to protect than storage systems in individual businesses. However, taking this step would require closer regulatory monitoring and oversight.

248. We acknowledge that there may be a case for a retailer's stores in a defined geographic area to share watchlists of repeat offenders if there is a clear problem that FRT would genuinely help to address (such as evidence that offenders move from supermarket to supermarket, or evidence of organised crime operations in multiple locations). Such shared watchlists would need to be carefully justified, however, and the seriousness of offending would be a key criterion, particularly if shared watchlists could deprive people of opportunities to obtain food or other essential items. Recent retail crime statistics suggest that 10% of retail crime offenders cause more than 60% of the harm. Shared watchlists that target those more serious recidivist offenders are likely to be more easily justified than sharing a store's whole watchlist.

249. OPC would want to work with retailers and Police on the design, oversight and governance of any more centralised FRT system, should one be proposed.

Impact of other criminal justice measures?

250. There may be future developments that could affect the supermarket security environment. For example, the effect of new developments in the retail sector such as the proposed citizen arrest powers may affect decisions about future or ongoing use of FRT in supermarket settings as this could affect the current risk environment and the mix of interventions that are needed.

9 Conclusion

251. Our Inquiry covers the specific privacy issues as they applied to FSNI's operating model for live FRT during the trial and has found that it complied with the Privacy Act.

252. It is also useful for others. While this report is not a green light for more general use of FRT, we recognise the importance of the issue for many businesses. What we have learned from this trial will enable other businesses to ask themselves the right questions about whether FRT is necessary and appropriate for them and will give them strong guidance about what they would need to do to set FRT up in a privacy-safe way.

253. As with any introduction of a technology or process that uses personal information, each new use of FRT must be considered on its own merits and be carefully justified in a business' individual context. It must assess the privacy impacts on members of the community, as well as making sure that FRT will be effective to help prevent crime in a particular setting. All elements of the FRT operating model – not merely the technology itself – must be well designed. This report shows businesses what those elements are and what they would need to do. It will not provide all the answers, but it will help.

254. To be clear, we would not expect every business to trial the technology in the comprehensive way that FSNI has done. Sometimes FRT will self-evidently be effective to meet a specific lawful purpose. Sometimes, the business environment will be clearly comparable to FSNI's situation and the findings from this trial will be enough to show that FRT would work for them. However, even if a business does not run a trial as such, we recommend testing its FRT settings pre-deployment and over time so it can make any necessary adjustments (for instance to its watchlists, camera placement or staff training) to ensure the system operates effectively and safely.

255. We will continue to watch this fast-developing area and will comment or provide further guidance if necessary. If we receive complaints or become aware of concerns, we will ask the agency to justify its FRT use settings and can take enforcement action if appropriate.

Appendix 1: Inquiry Terms of Reference

Privacy Commissioner Inquiry into Facial Recognition Technology Trial by Foodstuffs North Island, NZ

April 2024

The Privacy Commissioner is conducting an Inquiry under section 17(1)(i) of the Privacy Act 2020 into the Facial Recognition Technology (FRT) trial conducted by Foodstuffs North Island (FSNI) in selected supermarkets across their network.

Facial recognition technology uses algorithms to try to identify a person by scanning and creating a digital map of that person's face and comparing it to a database of facial images held by the organisation.

Biometric information, such as the scan of a person's face, is sensitive personal information and is regulated under the Privacy Act 2020.

Inquiry Rationale – Developments in FRT

Where facial recognition is used in New Zealand, it tends to be in situations where an individual is seeking to verify their identity to access a device or a service. The use of live facial recognition technology to scan and identify an individual in real time and compare them against a database of faces is rare in New Zealand.

FSNI is considering using live facial recognition technology as an additional tool to address retail crime in their supermarkets, with a primary focus on violent retail crime and keeping their staff and customers safe.

There is no known other current use of facial recognition technology in retail in New Zealand. Supermarkets are an essential service provider. As such there is high public interest in the operation and outcome of this trial.

FSNI Facial Recognition trial

FSNI is undertaking an independently evaluated trial of FRT in 25 selected franchised supermarkets over six months to provide evidence of effectiveness. The trial stores have been rolled out in four tranches with the final tranche going live on 8 March 2024.

During the trial, FRT will be used to scan and make a biometric face template of each customer as they enter trial supermarkets to see if they match a watchlist of people identified as causing harmful behaviour. FSNI have defined harmful behaviour as being theft, burglary, robbery, assault (physical or verbal) and other aggressive, violent, and threatening behaviour. FSNI has incorporated several privacy mitigations into their operational protocols, including on the storage and use of images.

Authority

This Inquiry is a Privacy Commissioner initiated Inquiry under section 17(1)(i) of the Privacy Act 2020. The Privacy Commissioner has Inquiry functions under this provision to inquire generally into any matter including any practice or procedure, whether governmental or non-governmental, or technical development, if it appears to the Privacy Commissioner that the privacy of the individual is being, or may be, infringed.

As part of the Inquiry, the Privacy Commissioner may use his powers to summon witnesses and obtain relevant information and documentation under section 203, referencing sections 86 through to 90 of the Privacy Act 2020. Information and documentation collected as part of this Inquiry will be held in confidence under section 206 of the Privacy Act 2020, the Office of the Privacy Commissioners obligation of secrecy, and is privileged under section 90 of the Privacy Act 2020.

Matters for Inquiry

The purpose of this Inquiry is to gather information to monitor the implementation of the FSNI trial to ensure it is compliant with the Privacy Act 2020, and to ensure that privacy risks are identified and appropriately mitigated as set out in FSNI's Privacy Impact Assessment. The Inquiry findings will inform the Commissioner's assessment of the level of any residual privacy risks, and the effectiveness of the use of FRT in reducing harmful behaviour in FSNI supermarkets.

Additionally, OPC is concerned about known bias and accuracy issues when it comes to FRT, particularly for people of colour. In the New Zealand context, our Office is particularly interested in the impact of FRT for Māori, Pasifika, Indian and Asian shoppers.

Scope/Methodology

The Inquiry will consider the following matters during the trial and evaluation period to ensure compliance with the Privacy Act and assess the impact of the trial on personal privacy:

1. Review the operational protocols set out by FSNI, including those in the user manual and Privacy Impact Assessment.
2. Review the on-the-ground implementation of these protocols in selected supermarkets, including:
    (a) The identification of individuals of interest for inclusion on FRT watchlists and processes for appealing inclusion and removal from watchlists, including audits of selected supermarket's watchlists
    (b) Transparency and notification practices
    (c) Information collection, use (including any secondary use), disclosure, retention and deletion
    (d) Staff training, including content and assessing staff understanding of the Privacy Act requirements
    (e) Access and correction and security settings for FRT information
    (f) Information flows between relevant parties, for example individual supermarkets and FSNI, and between FSNI and independent evaluator and the FRT software provider
    (g) Application and effectiveness of the mitigations to the privacy risks identified in the Privacy Impact Assessment
    (h) Number and type of alerts generated (both positive and false positive matches) and subsequent actions taken by supermarkets.
    (i) Number of instances where human checkers agree/don't agree with FRT match and the characteristics of the individual/context for each instance.
3. Key statistics across the trial where available, for individual trial stores and non-trial stores for comparative purposes. For trial stores this includes total number of faces scanned, total number of facial images deleted, numbers enrolled on watchlists, rationale for enrolment by specific category, action taken in response

to identification, referrals to police, subsequent police action. For non-trial stores comparative information includes estimates of total customer numbers, and comparative information from the Auror platform.

4. Review supermarkets watchlist and the number and details of adverse actions taken by staff to review any privacy impact of the trial on Māori, Pasifika, Indian and Asian shoppers given potential bias and accuracy issues associated with FRT.
5. Assess the findings of the trial evaluator at key points during the trial and at the completion of the evaluation report.
6. Seek and review feedback on the trial from customers, members of the public, selected supermarkets' owner operators, security staff, privacy officers, checkout operators.
7. Review complaints received by OPC, by individual FSNI supermarkets and FSNI to understand the nature and number of issues or concerns being raised by the public.
8. Evaluate the trial data and conclusions regarding the effectiveness of FRT in reducing the impact of retail crime and protecting staff and customers against harmful behaviour, including in comparison to other deterrence methods used.
9. Engage with Māori and Pasifika communities in the areas of the trial stores to identify any impacts particular to them.
10. Seek independent expert advice as needed.

At the completion of the Inquiry, the Commissioner will use the information and evidence obtained to produce a report which may include findings. Any breaches of the Privacy Act identified during the trial will be raised immediately with FSNI. At the completion of the trial, the Commissioner will consider the findings and determine any appropriate actions as set out in our Compliance and Regulatory Action framework[1].

Exclusions

Individual complaints

The Inquiry will not investigate any individual's complaint. However, the fact that the OPC is undertaking this Inquiry does not preclude an investigation under Part 5 of the Privacy Act at the request of an affected individual. Individuals who believe their privacy may have been interfered with and wishing to consider a complaint to the Office of the Privacy Commissioner can find more information on this process at www.privacy.org.nz.

Call for Information

People or organisations who believe that they have information that is relevant to this Inquiry should contact the Office of the Privacy Commissioner at FRTinquiry@privacy.org.nz

Commencement of work

The final tranche of FSNI FRT trial stores was rolled out on 8 March 2024. The Inquiry will commence on 3 April 2024. At the completion of the trial, a draft report will be provided to FSNI. The Commissioner will release his findings once any comments made by FSNI have been considered.

[1] https://www.privacy.org.nz/about-us/what-we-do/caraf/

Appendix 2: Inquiry timeline

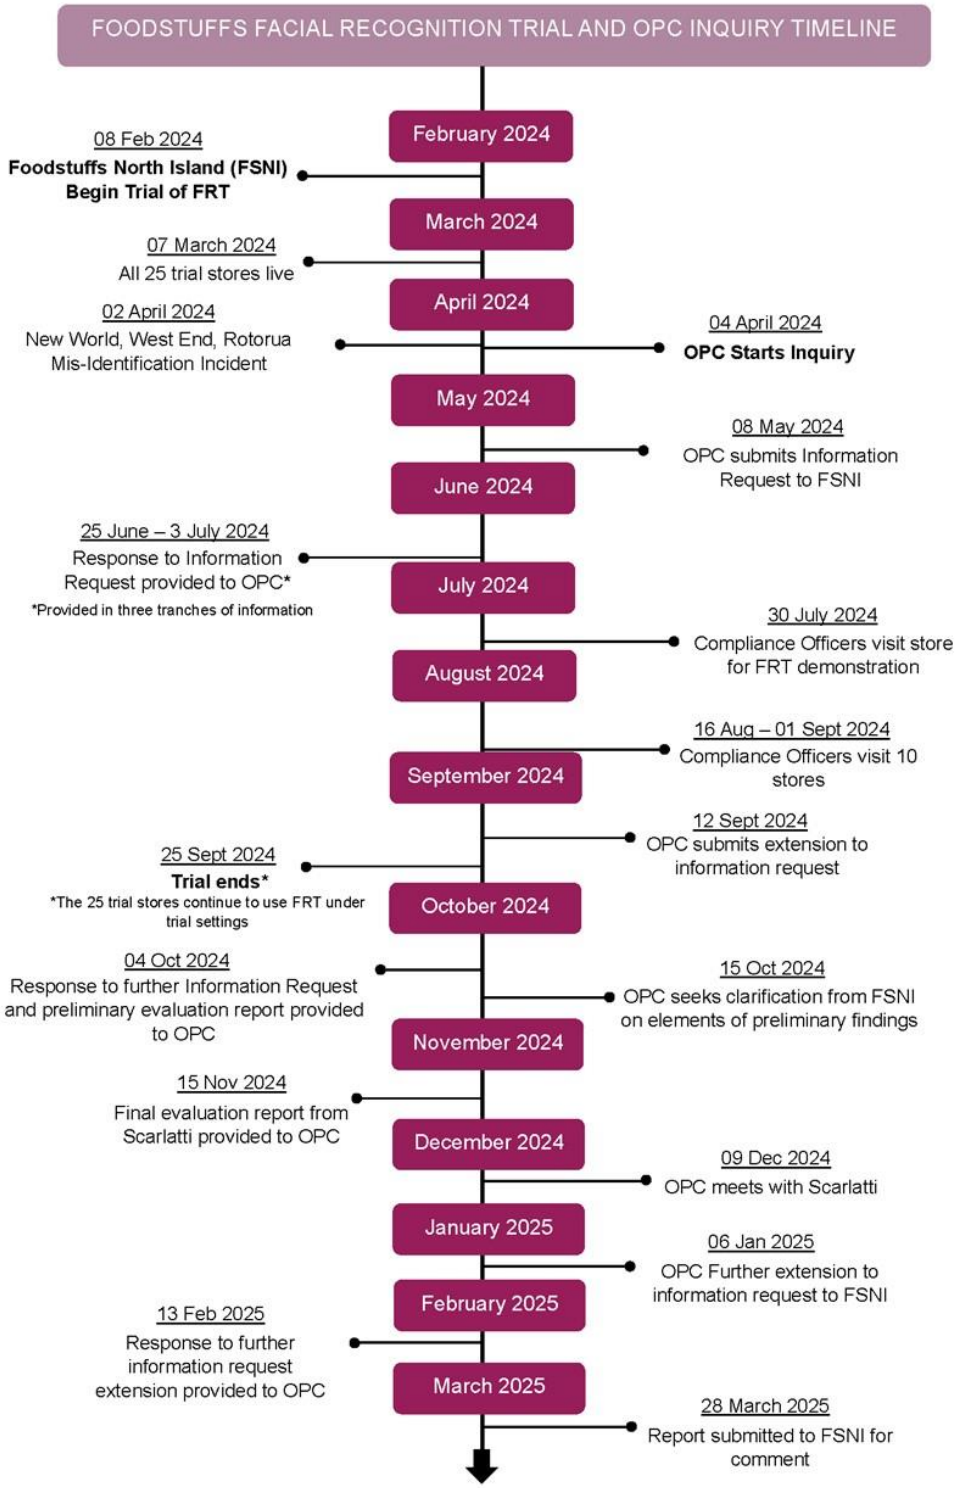

Appendix 3: List of participating and visited stores

Full list of stores participating in the trial	List of stores visited by Compliance Officers
PAK'n'SAVE Manukau, Auckland	PAK'n'SAVE Manukau, Auckland
PAK'n'SAVE Mill Street, Hamilton	PAK'n'SAVE New Plymouth
PAK'n'SAVE Napier	PAK'n'SAVE Ormiston, Auckland
PAK'n'SAVE New Plymouth	PAK'n'SAVE Palmerston North
PAK'n'SAVE Ormiston, Auckland	PAK'n'SAVE Sylvia Park, Auckland
PAK'n'SAVE Papakura, Auckland	PAK'n'SAVE Upper Hutt
PAK'n'SAVE Papamoa	PAK'n'SAVE Whanganui
PAK'n'SAVE Silverdale, Auckland	PAK'n'SAVE Whitiora, Hamilton
PAK'n'SAVE Sylvia Park, Auckland	New World Hillcrest, Rotorua
PAK'n'SAVE Taupō	New World Levin
PAK'n'SAVE Tamatea, Napier	
PAK'n'SAVE Upper Hutt	
PAK'n'SAVE Whanganui	
PAK'n'SAVE Whangārei	
PAK'n'SAVE Whitiora, Hamilton	
New World Brookfield Tauranga	
New World Feilding	
New World Gate Pa, Tauranga	
New World Hillcrest, Rotorua	
New World Kaikohe	
New World Levin	
New World Mt Maunganui	
New World Pioneer, Palmerston North	
New World Regent, Whangārei	
New World Te Rapa, Hamilton	

Appendix 4: Watchlist and alert flows

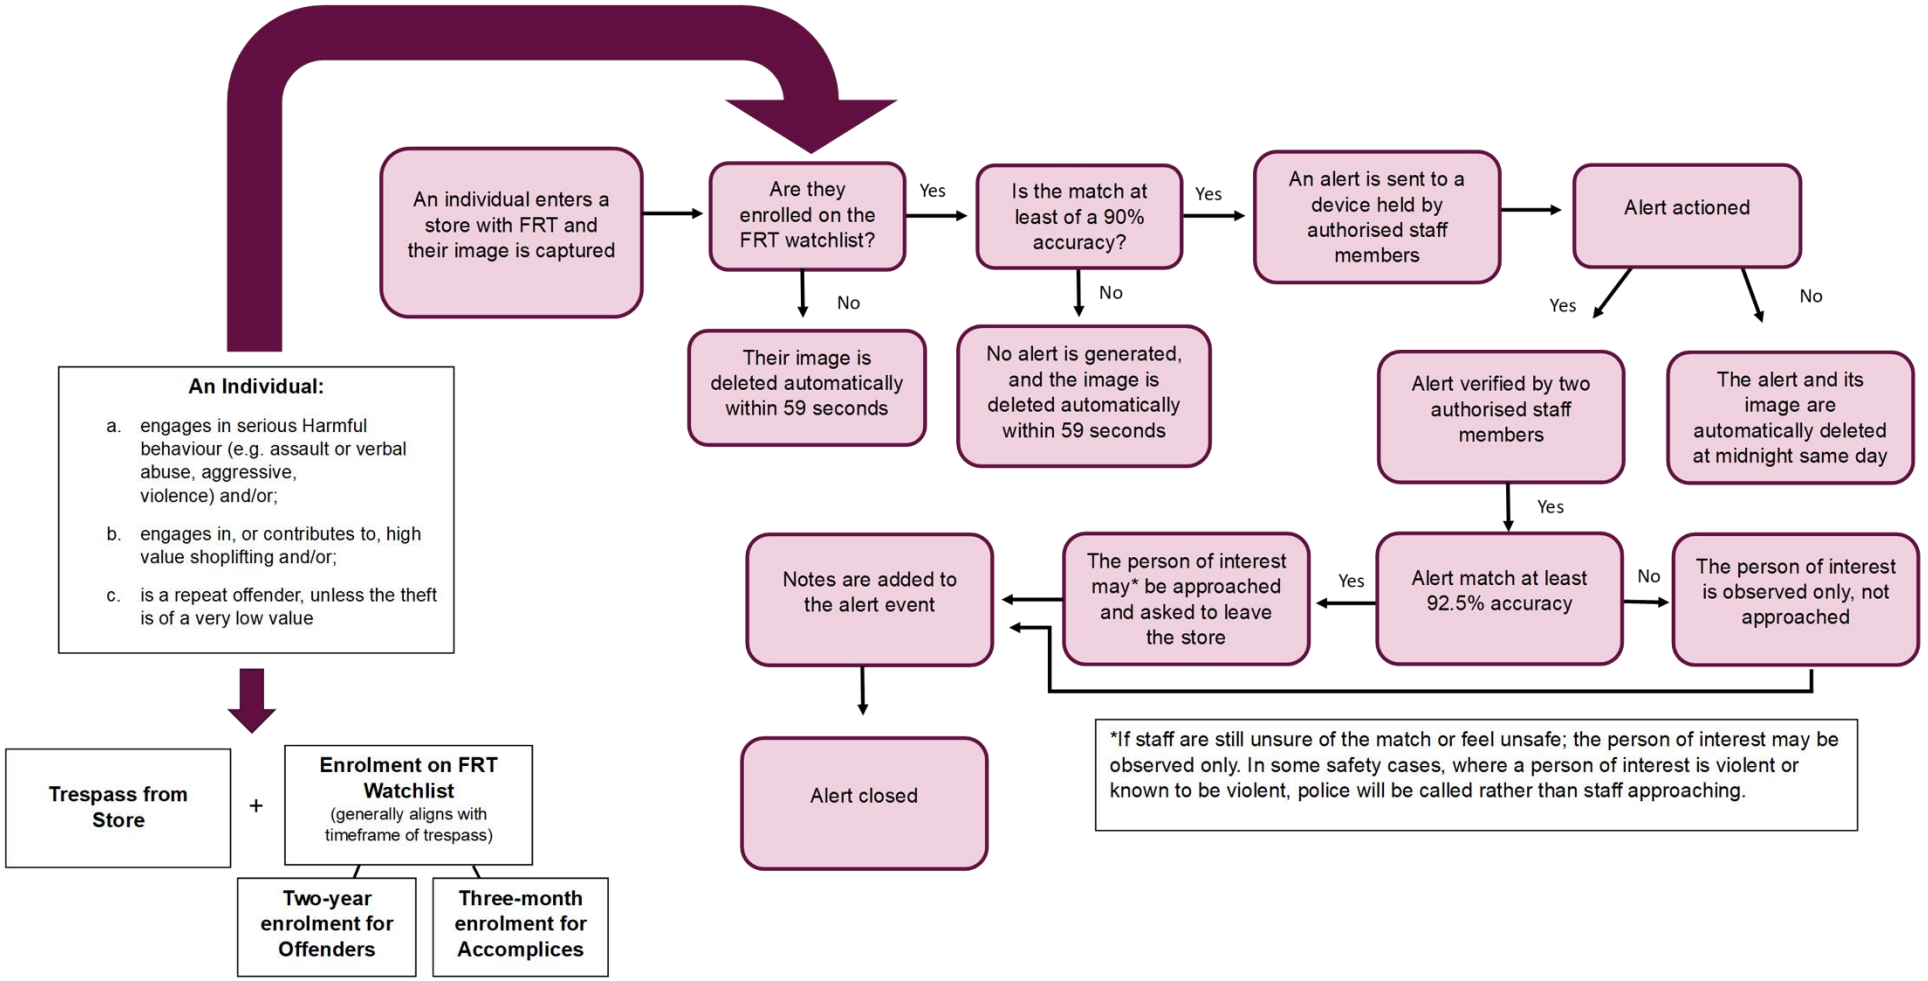

Appendix 5: The trial by the numbers