one.nz

AI Trust Report 2026

The Trust Imperative

Can New Zealand scale AI without losing public confidence?

AI Trust Report 2026

Foreword

Artificial intelligence (AI) is not a future possibility for New Zealand. It is already here, already visible, and already shaping how people interact with businesses, services and institutions. Three in four New Zealanders say they have interacted with an AI-powered tool or service in the past 12 months. But this year's research makes something else just as clear: building public trust is critical for New Zealand businesses and organisations in our journey to innovating with AI – particularly as we're operating in a global landscape where other countries are pushing ahead.

That matters, because the next phase of AI adoption will not be won by the organisations that move fastest alone. It will be won by the organisations that can show New Zealanders they are deploying AI responsibly, transparently, and with clear human accountability. The public is developing clearer expectations for AI. Kiwis still see value, especially in productivity and convenience, but they are increasingly asking harder questions about privacy, misinformation, fairness, oversight and control.

For Aotearoa New Zealand, this is a pivotal moment. AI could help lift productivity, improve services and support better national outcomes. But those gains will not be automatic. They will depend on whether people feel AI is being used in ways that are fair, useful, explainable and safe. We need to help build AI literacy at scale to support AI acceptance, especially among those who are currently lagging in terms of adoption and understanding. Otherwise, the digital divide will likely widen.

AI is also changing the economics of innovation itself. As the cost of building, testing and scaling new ideas rapidly falls, New Zealand has an opportunity to compete far beyond its size – but only if we build the trust, capability and confidence to move with it.

At One NZ, we believe the challenge is not whether to engage with AI. It is how to help shape an AI future that New Zealanders can trust – one that supports a better-connected future, with people still at the centre. Or what we define as taking an approach of "AI first, human where it matters most".

In the next phase of AI adoption, trust may depend less on whether AI is used at all, and more on whether AI systems can clearly prove who they are acting for, what authority they have, and how decisions can be traced and challenged. As agentic AI becomes more autonomous, identity, delegated authority and auditability are likely to become foundational trust requirements.

This report is intended to contribute to the national conversation, focusing on five key AI signals for Kiwi organisations. These signals indicate the key trust levers are deepening in importance, as the questions around AI adoption are changing. Leaders in business, government and policy now need to change with them.

Summer Collins
Chief AI & Business Services Officer, One NZ

"As AI moves from assisting people to acting on their behalf, trust becomes essential infrastructure for Kiwi organisations."

Contents

Executive Summary

62%

of Kiwis say they would stop using a product or service if they have concerns about how the organisation was using AI.

70%

of AI users say they have experienced at least one issue with AI-powered services in the past year.

68%

of New Zealanders say having a clear option to talk to a human if needed would make them more comfortable with having AI-powered customer interactions.

45%

of Kiwis are concerned about the environmental impacts of business AI use.

The five signals for New Zealand organisations

1. AI is becoming mainstream, but meaningful adoption remains shallow
2. Public sentiment is shifting from curiosity to scrutiny
3. Trust is being shaped by experience, not perception
4. Human oversight is becoming the baseline for trust
5. AI's value case is narrowing as expectations of guardrails rise

Overview of Insights

New Zealand has entered a new phase in the AI conversation. This second annual AI Trust Report and research shows that the defining tension in 2026 is not low AI awareness among the general public. It is the widening gap between broad exposure and fragile trust. AI is becoming normalised in everyday life, but confidence in how organisations use it has weakened over the past year. Public sentiment in Aotearoa is shifting from curiosity to scrutiny. The question is no longer simply whether AI is useful. It is whether it is being used by businesses and organisations in ways people believe are appropriate, accountable and worth the trade-offs.

The 2026 One NZ AI Trust Report shows that 76% of New Zealanders have interacted with an AI-powered tool or service in the past 12 months, stable compared to 2025 research, yet meaningful adoption remains concentrated in a relatively small active core. However, trust in organisations to use AI ethically and responsibly has declined across business types and sectors.

Fewer than two in five Kiwis now believe AI has the potential to deliver better outcomes for society, and 70% of AI users say they have experienced at least one issue with AI-powered services in the past year. The decline in trust is most pronounced among the most engaged users, suggesting this is being shaped by direct experience, not distant suspicion.

This shift has practical consequences. New Zealanders are still open to AI's value, but that value is narrowing. People see the case for AI in reducing admin, lowering operational costs and speeding up customer service. They are less convinced by broader promises. At the same time, expectations of guardrails are rising. Human oversight has become a more important trust condition. Clearer government standards would make nearly half the country (47%) more likely to trust AI-powered products and services. Sustainability sits squarely within the trust equation as well, with 45% concerned about the environmental impacts of business AI use.

For organisations in New Zealand, that points to a new strategic imperative. The competitive question is no longer just how quickly AI can be deployed. It is whether it can be deployed in a way that earns confidence. For government, the implication is just as sharp. If New Zealand wants AI to support productivity and long-term national outcomes, it needs public trust, clear expectations, and a credible framework for accountability.

The five signals for New Zealand organisations

1. AI is becoming mainstream, but meaningful adoption remains shallow.
2. Public sentiment is shifting from curiosity to scrutiny.
3. Trust is being shaped by experience, not perception.
4. Human oversight is becoming the baseline for trust.
5. AI's value case is narrowing as expectations of guardrails rise.

INSIGHT 1 AI is becoming mainstream, but meaningful adoption remains shallow

The first signal relates to adoption, or perceived use of AI, tools by Kiwis in their everyday lives.

AI is familiar to many New Zealanders. Net interaction with AI-powered services, or the understanding that they are interacting with an AI-powered service, has held steady at 76% over the past year. Younger New Zealanders continue to knowingly use AI far more than older groups, with interaction at 89% among 18–34s versus 62% among those aged 55+.

But the pattern beneath that headline is more revealing than the headline itself. AI interaction is still mostly occasional across categories, and only chat-based assistants show meaningful levels of regular use with 50% of Kiwis say they have used a chat-based AI assistant in the past year, such as ChatGPT, Perplexity, Claude or Gemini. In other words, AI is visible, but it has not yet become deeply embedded in everyday behaviour for most people. This creates a risk for leaders: mistaking visibility for readiness.

The adoption gap matters. The research shows a relatively small AI Active segment, representing 29% of the population, alongside a much larger passive or avoidant majority. AI Actives are the pioneers: they use AI more often, feel more confident with it, and are more optimistic about its potential to improve customer and societal outcomes. But even they are not blindly enthusiastic; they are just as likely to stop using a product or service if they have concerns about how AI is being used. AI Dabblers are the persuadable middle: open to experimenting, but still looking for reassurance, confidence and proof from others before leaning further in. AI Avoiders are the hold-outs: less confident, less trusting, and more likely to be driven by lack of knowledge, general distrust, and concern about what AI may replace.

AI Engagement
(In the past 12 months)

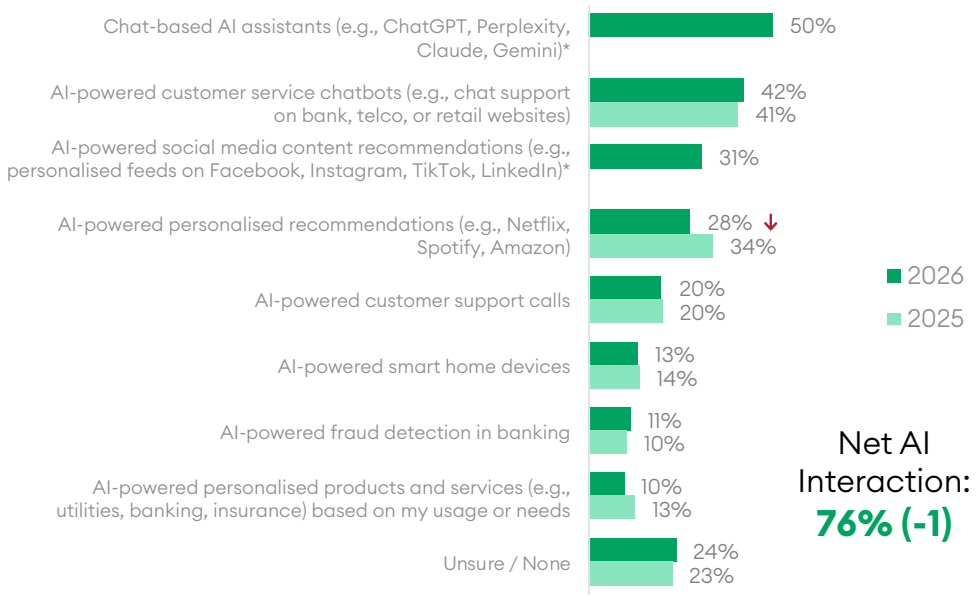

Net AI Interaction:
76% (-1)

* New option added 2026
Q. To the best of your knowledge, have you interacted with any of the following AI powered services from businesses / organisations within the past 12 months? Please select all that apply.
Base: Total Sample 2025 (n=1,000); 2026 (n=1001)

Significantly higher / lower at 95% confidence ↑↓ | *+/- Comparisons are against 2025

INSIGHT 1 AI is becoming mainstream, but meaningful adoption remains shallow

This is an important strategic warning. Exposure should not be mistaken for acceptance. It is easy for leaders to see high levels of interaction and conclude that the public is ready to go further. The report suggests otherwise. What New Zealand has today is broad awareness, shallow adoption and highly uneven readiness. The organisations that misread that gap will over-automate too early, and then wonder why public resistance hardens.

There is also a wider national implication sitting underneath this signal. If younger New Zealanders are already engaging with AI at materially higher rates than older groups, then the next stage of AI deployment risks widening capability gaps rather than narrowing them. If AI becomes embedded in how essential services, advice, healthcare, finance or customer support are delivered, then accessibility and confidence are no longer side issues – they are central to whether AI improves outcomes at all.

This points to a deeper shift. AI literacy is quickly becoming a form of economic infrastructure, not just a technical skill. As AI becomes embedded in how work is performed and services are delivered, differences in confidence and capability are likely to translate into differences in productivity, opportunity and outcomes. For organisations, workforce confidence will directly shape how far AI can be deployed.

For the country, uneven capability risks reinforcing existing inequalities rather than reducing them. Trust is also closely tied to understanding – people are more likely to engage with systems they feel they can interpret, question and control.

AI Engagement Segments

AI ENGAGEMENT: LOW → HIGH

	AI Avoiders 24% (+1)	AI Dabblers 48% (-4)	AI Actives 29% (+3)
OVERVIEW	Little to no recent interaction with AI Low trust and low perceived relevance Least convinced of broader benefits	Have tried AI, but usage is irregular and shallow Lack confidence which limits progression to habitual use Less likely to see value outside of more simple use cases	Confident, frequent users across personal + professional contexts Over-index on belief in AI's societal value Early adopters and influencers in their networks
SKEW	**Older** (38% among 55+ vs 11% among 18 - 34) **Regional** (28% vs 21% among those in main centres) **Lower Income earners** (28% among <$100K Household Incomes vs 15% among $100K+)	No significant demographic skews - **straddle a broad range of demographics.**	**Younger** (39% among 18 - 34-year - olds vs 14% among 55+) **Major centers** (32% vs 24% among regional Kiwis) **Mid/higher income earners** (39% among $100K+ Household Incomes vs 22% among <$100K).
IMPLICATIONS	*Conversion will be slow – significant trust barriers to overcome.*	*Largest segment – key opportunity for organisations to shift perceptions.*	*Small group – outsized influence on adoption and advocacy*

*+/- Comparisons are against 2025

INSIGHT 1 AI is becoming mainstream, but meaningful adoption remains shallow

AI Actives are frequently using AI in their day to day lives and leading the way when it comes to adoption of new tools & services.

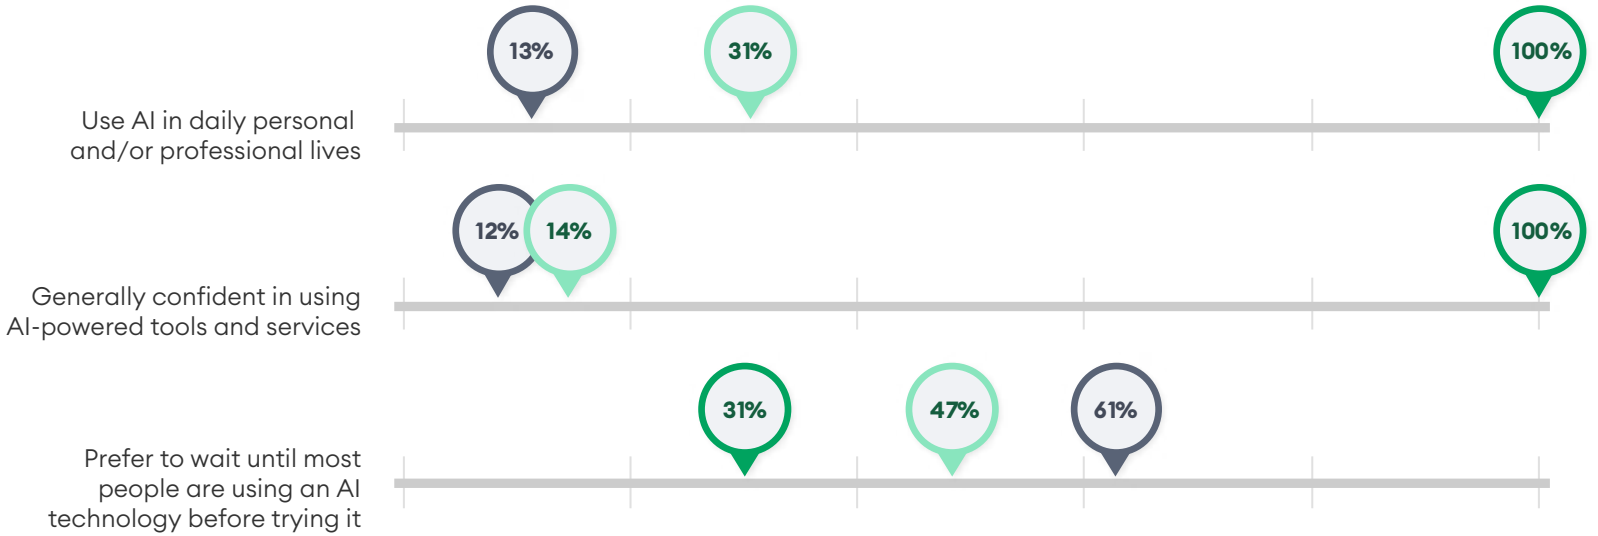

Q. To what extent do you either agree or disagree, with each of the following statements related to AI?
Base: Total Sample 2025 (n=1,000); 2026 (n=1001)

INSIGHT 1 AI is becoming mainstream, but meaningful adoption remains shallow

AI Actives also share greater optimism around AI's societal impact, however, are as likely to stop using a product or service if they have concerns about how AI is being used.

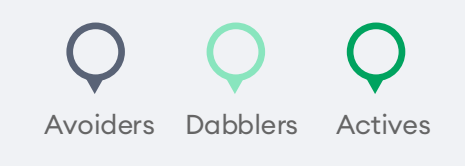

Q. To what extent do you either agree or disagree, with each of the following statements related to AI?
Base: Total Sample 2025 (n=1,000); 2026 (n=1001)

INSIGHT 2 Public sentiment is shifting from curiosity to scrutiny

The second signal points to a shift in sentiment and mood among New Zealanders towards how organisations are using AI.

Consistent with 2025, almost two-thirds (62%) of Kiwis say they would stop using a product or service if they have concerns about how the organisation was using AI. However, the most important year-on-year change in the report is not about usage. It is about mood. AI usage has held up, but trust and comfort have declined.

Trust in organisations to use AI ethically and responsibly has fallen across sectors and organisation types. Multinationals remain the least trusted, with 51% of New Zealanders expressing low or no trust in their responsible use of AI. Small businesses and large New Zealand businesses have also seen some of the sharpest drops. At the same time, belief that AI will deliver better outcomes for society has weakened, falling from 43% in 2025 to fewer than two in five in 2026.

This is the real meaning of the shift from curiosity to scrutiny. The public is moving beyond AI as a novelty. New Zealanders are now asking tougher questions: Is this accurate? Is it fair? Is it safe? Is my data protected? Who is accountable if something goes wrong? These are not fringe concerns. They are becoming the mainstream test for whether AI feels acceptable in practice. The rise of misinformation and deepfakes, combined with growing concern about data use and weak controls, has made people more suspicious of the environment AI is operating in.

Levels of Trust
(That AI would be used ethically and responsibly)

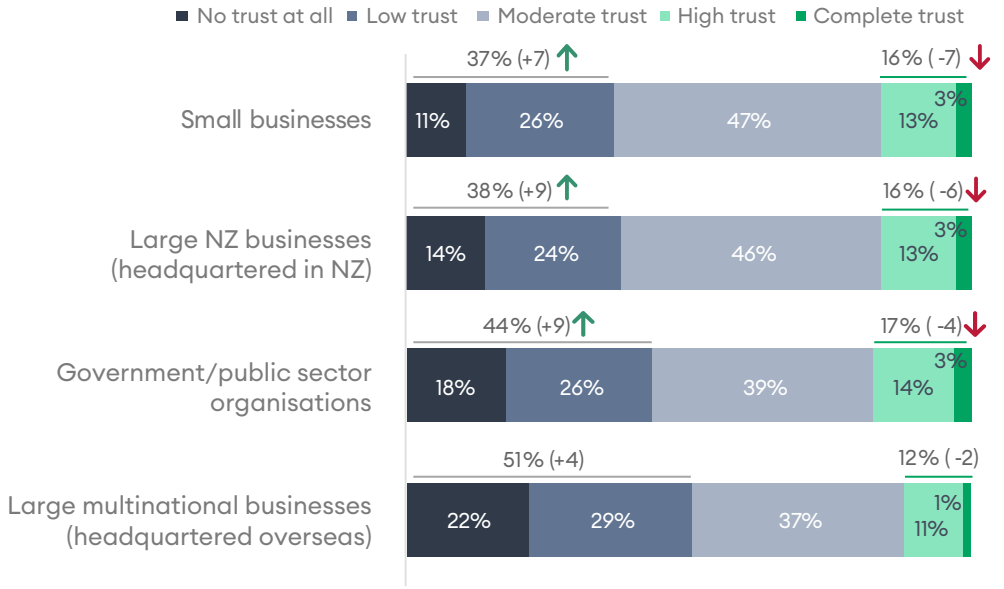

Q. To what extent do you trust the following types of businesses / organisations to use AI ethically and responsibly? i.e. ensuring fairness, transparency, privacy, and accountability for their customers
Base: Total Sample 2025 (n=1000); 2026 (n=1,001)

Significantly higher / lower at 95% confidence ↑↓ | *+/- Comparisons are against 2025

INSIGHT 2 Public sentiment is shifting from curiosity to scrutiny

The industry lens sharpens the point. New Zealanders remain most comfortable with AI in retail and telecoms, but even there comfort has declined. In sectors where stakes are higher, comfort is much lower: only around a quarter are now comfortable with AI use in healthcare, and under a third in finance and banking. Government sits in the same zone of caution. This matters because these are exactly the sectors where AI could have high public impact – and exactly the sectors where public legitimacy will be hardest won.

The implication for leaders is more confronting than it may appear. Many organisations are still talking about AI as if the country is in an early adoption phase defined by excitement and experimentation. The report suggests that phase is ending. New Zealand is entering a judgement phase. People are no longer asking only what AI can do. They are asking where it belongs, where it does not, and what protections need to come first.

Comfort With AI Use by Industry

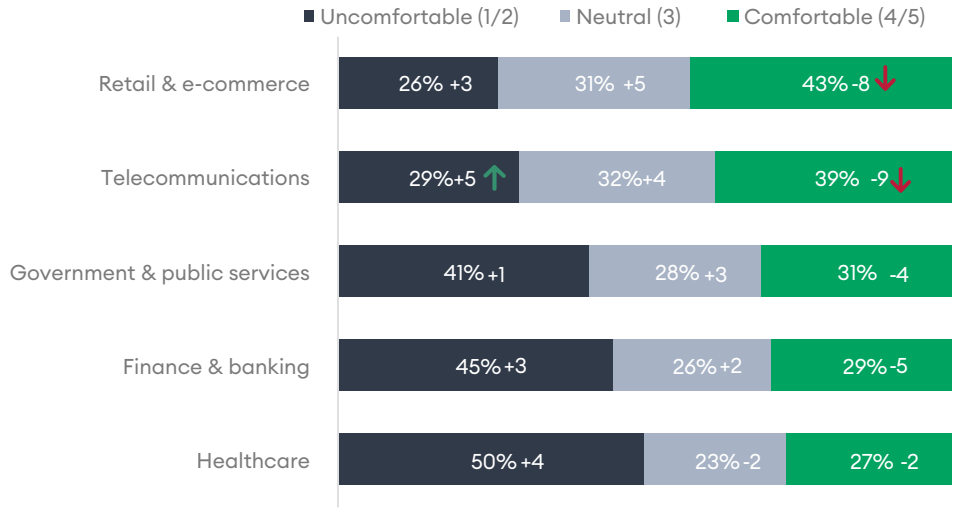

Q. How comfortable would you be using AI if you were dealing with an organisation in the following industries?
Base: Total Sample 2025 (n=1,000); 2026 (n=1001)

Significantly higher / lower at 95% confidence | *+/- Comparisons are against 2025

INSIGHT 3 Trust is being shaped by experience, not perception

The third signal Kiwi organisations should be paying attention to centres on how user experiences of AI are shaping trust.

One of the most revealing findings in the report is that trust is declining fastest among those with the most direct experience of AI. Regular users remain more trusting overall than non-users, but they have seen the sharpest drop in trust over the past year for all types of organisations using AI. That matters because it challenges a common assumption: that trust will naturally rise with familiarity. In practice, the opposite may happen when real experiences do not match the promise. AI does not behave like traditional technology tools, and variability in results is becoming more visible, particularly for AI Actives.

The numbers behind that are hard to ignore. Seven in ten AI users say they have experienced at least one issue with AI-powered services in the past 12 months, with two issues experienced on average. The most damaging problems are not theoretical. They are immediate and practical: incorrect information, poor outcomes, privacy concerns, difficulty reaching a human, and experiences that feel impersonal or frustrating. Younger users are more likely to report wrong answers and security concerns. Older users are more likely to report the loss of human connection – not being able to reach a person when needed, or feeling the experience was cold and frustrating. These experiences are not neutral; they materially reduce people's willingness to use similar AI-powered services again.

High / Complete Trust by User Groups

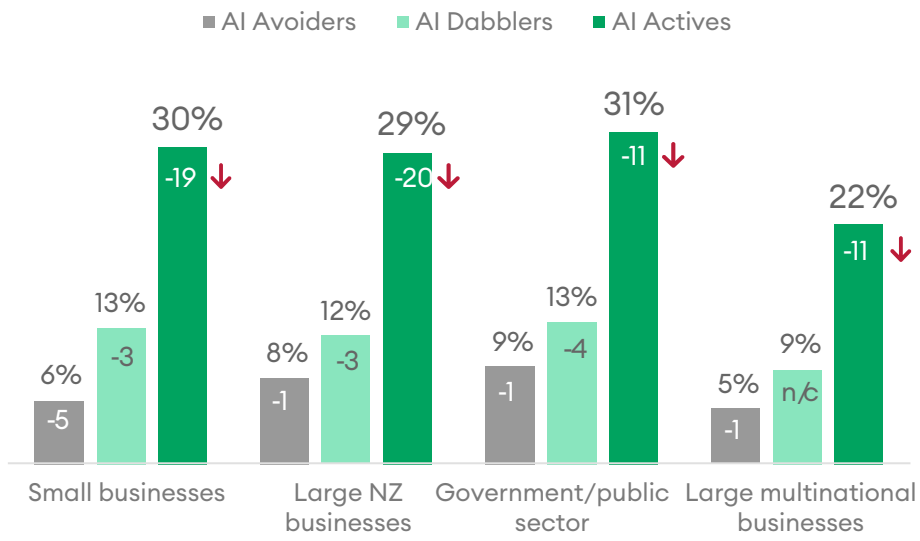

Q. To what extent do you trust the following types of businesses / organisations to use AI ethically and responsibly? i.e. ensuring fairness, transparency, privacy, and accountability for their customers
Base: Total Sample 2026

AI Actives (n=289), AI Dabblers (n=478), AI Avoiders (n=238)

Significantly higher / lower at 95% confidence ↑↓ | *+/- Comparisons are against 2025

INSIGHT 3 Trust is being shaped by experience, not perception

In practice, this creates a form of trust asymmetry. A single instance of incorrect information, a privacy concern or a poor experience can have a disproportionate impact on how people view not just that service, but AI more broadly. Individual deployments do not operate in isolation – they shape expectations for the wider ecosystem.

This is why governance maturity matters early. Poor implementations do not just create local issues; they can slow adoption more broadly by reinforcing scepticism. Trust becomes harder to rebuild than to lose. This should not be treated as a communications problem. It is a performance and operational problem. The public is not simply nervous because AI is unfamiliar. Many are becoming sceptical because they have already seen where it falls short – whether through unreliable outputs, unclear data practices, or a perceived lack of accountability when things go wrong.

This points to a deeper shift in how trust must be managed. Trust in AI is no longer sustained by policy statements or ethical principles alone. It increasingly depends on an organisation's ability to operate AI systems reliably over time – to monitor performance, detect issues early, respond when things go wrong, and provide clear accountability for outcomes.

Issues with AI-powered services
(In the past 12 months – Among those who have interacted with AI services)

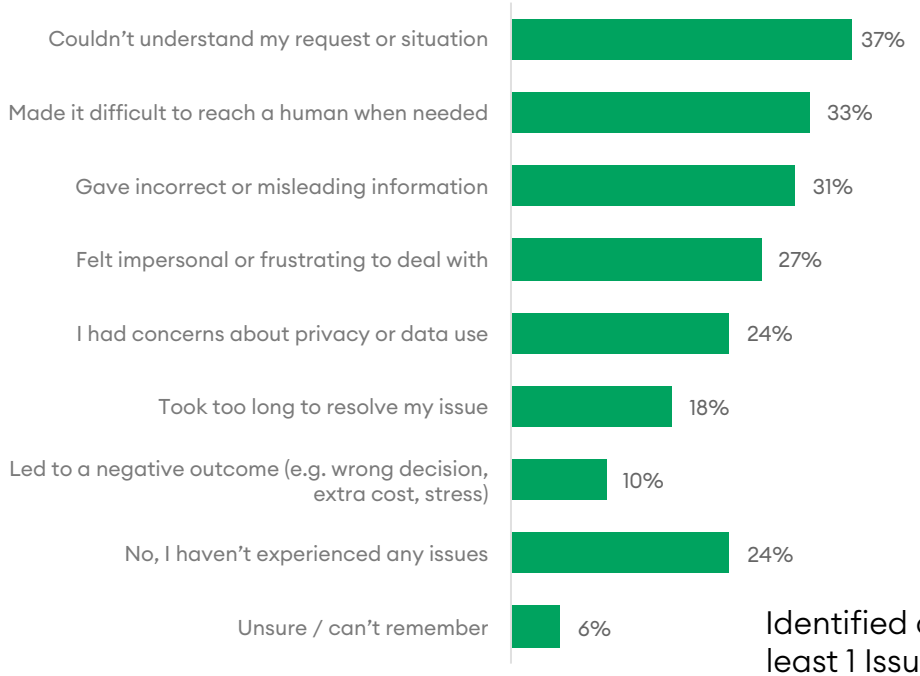

Identified at least 1 Issue:
70%
(2 issues experienced on average)

**New Question added in 2026

Q. Have you experienced any issues with the AI powered services you've used in the past 12 months? Please select all that apply
Base: 2026 Those who have interacted with any AI powered services (n=735)

INSIGHT 3 Trust is being shaped by experience, not perception

This highlights an important dynamic in how trust in AI is formed. Trust is not built and lost evenly. It tends to accumulate gradually through repeated positive experiences, but can be eroded quickly by a small number of negative ones.

In that sense, trust is becoming an operational discipline. It requires ongoing oversight, traceability of decisions, effective escalation pathways, and the ability to learn from and correct failures quickly. As AI systems become more embedded in customer experiences and decision-making, these capabilities begin to resemble other core operational functions, such as cybersecurity or financial controls.

The implication is blunt. Organisations that deploy AI badly are not just creating isolated customer service issues – they are teaching the public to trust AI less. In that environment, every poor deployment has a wider effect, shaping expectations for the next system, the next service, and the next organisation.

Impact of Negative Experience

(Made me much less likely to use this kind of AI-powered service / stopped using altogether)

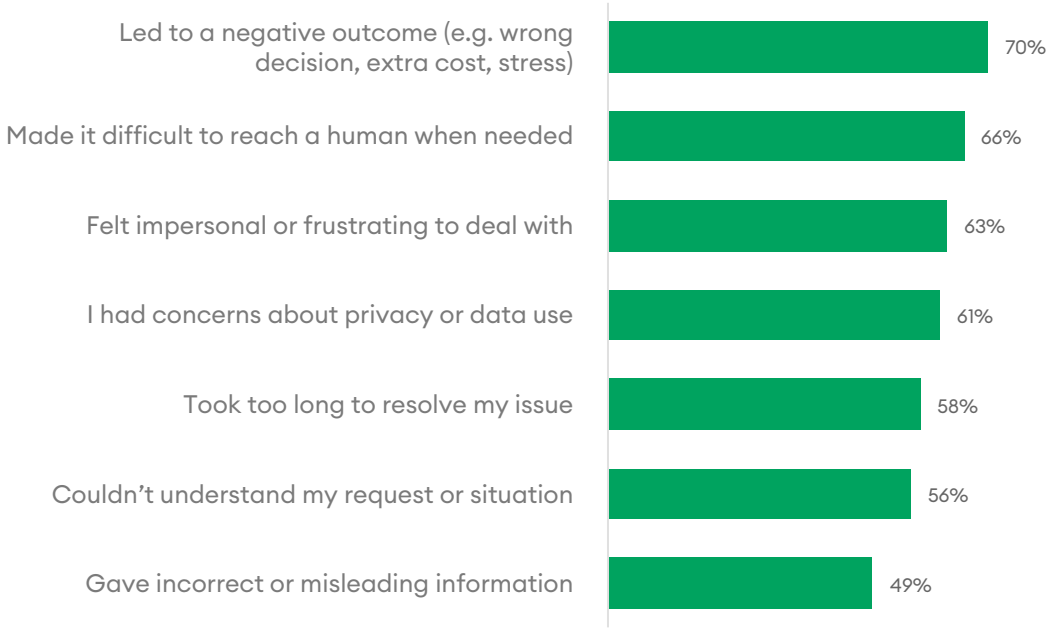

**New Question added in 2026

Q. You mentioned you've experienced the issues below, when using AI powered services in the past 12 months. What impact have those experiences had on your level of AI use?
Base: Those who have experienced issues with AI - powered services 2026 (n=504)

INSIGHT 4 Human oversight is becoming the baseline for trust

Signal four relates to how to build trust. Humans like to know another human is in charge of AI-powered customer interactions.

As trust erodes, New Zealanders are not uniformly calling for less AI. They are calling for more visible human accountability around it. In the research, the only action that has increased in importance as a trust-builder is evidence of human oversight. That signal is particularly strong among AI Dabblers, with 63% saying this would make them more comfortable with AI-powered customer interactions. At the same time, 32% of AI Avoiders say there is nothing organisations could do to make them more comfortable, suggesting that some portion of scepticism may be difficult to reverse once it sets in.

The persuadable middle, however, is giving leaders a very clear instruction. People want to know that a human remains in the loop, or at the helm, especially when stakes are high. Healthcare, finance and government are the industries where evidence of human oversight matters most – and they are also the sectors where comfort with AI is lowest. That is not a coincidence. It shows that public trust is not only about the technology itself, but whether people believe they can still access judgment, escalation and accountability when it matters.

The expectation is not simply that humans supervise automated systems, but that they remain responsible for higher-order judgment, context and relationship where it matters most. In practice, AI is likely to take on more routine and repeatable tasks, while human roles shift toward interpretation, exception handling, empathy and decision-making in complex or ambiguous situations. The future is not a choice between humans or AI, but a rebalancing. Trust is shaped not only by whether a human is available as a fallback, but by whether people feel that human judgment, accountability and connection are still present where they matter most.

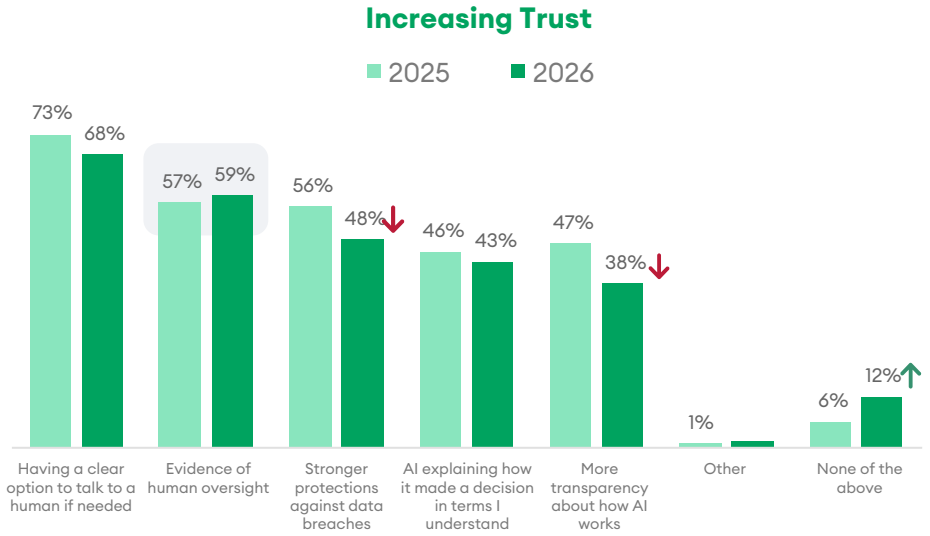

Q. Which of the following would make you more comfortable with having AI-powered customer interactions with businesses in future? Please select all that apply.
Base: Total Sample 2025 (n=1,000); 2026 (n=1001)

Significantly higher / lower at 95% confidence | *+/- Comparisons are against 2025

INSIGHT 4 Human oversight is becoming the baseline for trust

This finding deserves to be taken literally. Human oversight cannot just exist in policy – it has to be legible in the experience. A buried disclosure statement is not the same as visible accountability. A chatbot that technically allows escalation, but makes it difficult to reach a person, does not feel like human oversight in practice. Organisations need to design for the moment of doubt: when a customer wants clarification, challenge, exception handling or human reassurance. If that moment is poorly handled, trust falls sharply.

The implication is that the trust equation is becoming clearer, not fuzzier. Data privacy, regulatory compliance and human oversight remain the baseline conditions New Zealanders expect. The question now is whether organisations are prepared to deliver them in practice, not just in principle.

Even this is unlikely to be the end state. A further shift is already emerging beneath the surface of this research. Today, most AI interactions are still framed as tools, assistants or systems that support human decision-making. But the next phase of AI adoption is likely to involve systems that act with greater autonomy – carrying out tasks, making decisions within defined parameters, and in some cases operating on behalf of customers or organisations.

This shift materially changes the trust equation. When AI moves from advising to acting, expectations go beyond human oversight alone. Questions of identity, delegated authority and permission become more important: who or what the system is acting for, what authority has been granted, and how those actions can be traced, challenged or reversed.

While this report reflects current public experience with AI, it also signals that trust frameworks will need to evolve quickly. The conditions New Zealanders are already calling for – transparency, accountability and human oversight – are likely to become baseline requirements in a more autonomous AI environment.

Where are Human Interactions Most Important?

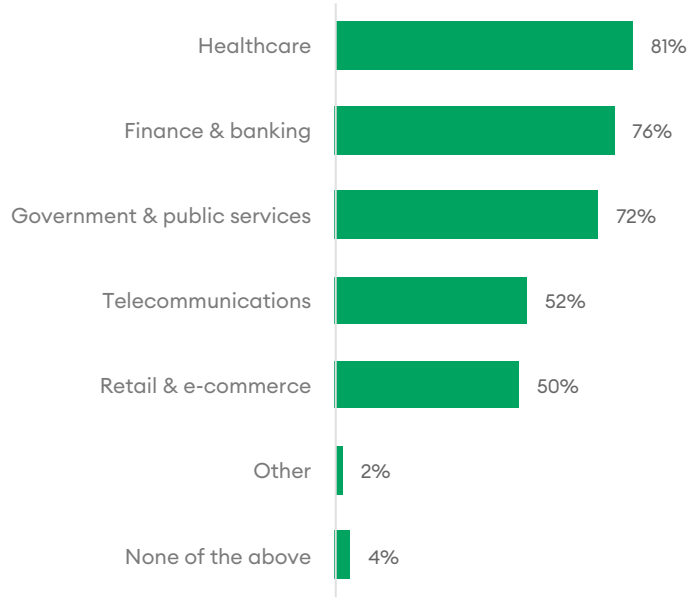

**New Question Added 2026

Q. You mentioned that you would be more comfortable having AI-powered customer interactions with businesses in future if you had evidence of human oversight. In which industries would this be important? Please select all that apply.**
Base: Those who would be more comfortable with human oversight (n=601)

INSIGHT 5 AI's value case is narrowing as expectations of guardrails rise

The fifth and final signal relates more to responsible AI, and the benefits New Zealanders expect from this emerging technology.

The report does not suggest New Zealanders have turned against AI outright. But it does show that the public case for AI is tightening. The strongest perceived benefit remains productivity, cited by 41% of respondents. By contrast, confidence in broader benefits has softened. Perceptions of AI improving cybersecurity and fraud prevention have dropped to 34%, down eight points year on year. Perceived benefit to public infrastructure has fallen to 25%, down six. Nearly one in five New Zealanders cannot name any benefit of AI at all, rising to 28% among those aged 55 and over.

Areas AI can Positively Impact NZ

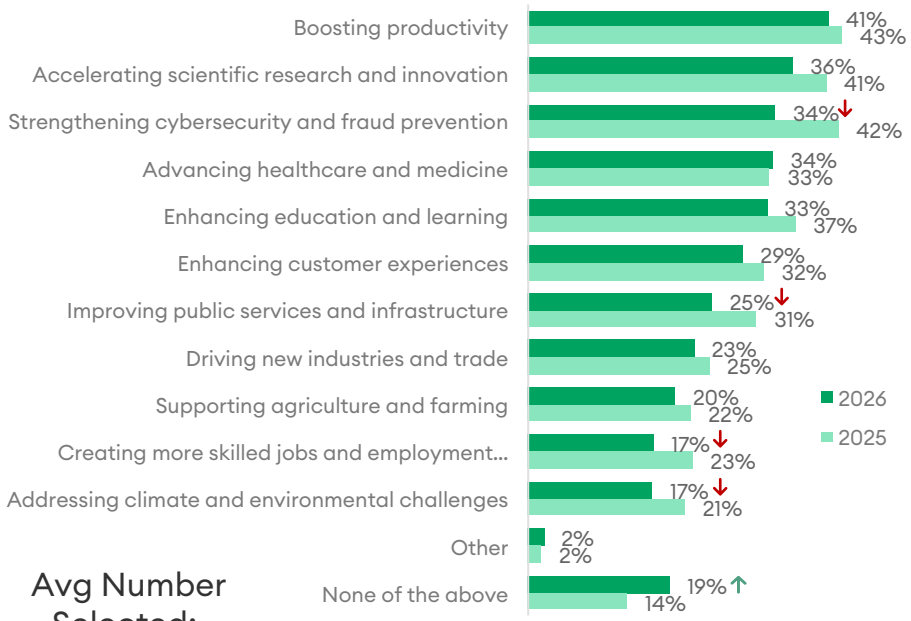

Avg Number Selected: **3.1 (-0.4)**

Q. In which areas do you think AI can have the biggest positive impact on New Zealand? Please select all that apply.
Base: Total Sample 2025 (n=1,000); 2026 (n=1,001)

Significantly higher / lower at 95% confidence ↑↓ | *+/- Comparisons are against 2025

INSIGHT 5 AI's value case is narrowing as expectations of guardrails rise

That tells its own story. The public is still open to AI, but mostly when the value is practical, immediate and believable. When people do see productivity gains, they are grounded in reducing admin, lowering operational costs and speeding up customer service. Expectations are much lower when it comes to new products, new value creation or broad social transformation. This is a more disciplined, less romantic public mood. It says: show me the concrete benefit first. Do not start with the grand promise.

Beyond this, there is a deeper structural shift beginning to emerge. While much of the current conversation focuses on AI improving efficiency within existing processes, the longer-term impact is likely to be on how work itself is designed. As AI systems take on more routine, analytical and coordination tasks, organisations may begin to redesign workflows, roles and decision-making structures rather than simply automating individual tasks.

For New Zealand, this has particular significance. With a high proportion of small businesses, constrained labour pools and ongoing productivity challenges, AI has the potential to enable leaner, more flexible operating models – where smaller teams can deliver greater output, and where geographic distance becomes less of a constraint. Realising that potential, however, will depend on whether organisations are prepared to rethink how work is structured, not just how technology is deployed.

The same discipline appears in the public's expectations of how AI should be governed. Community wellbeing and efficiency gains are the top priorities people want organisations to focus on when using AI. Nearly half of New Zealanders (47%) say they would be more likely to trust AI-powered products and services if government set clearer expectations and standards. The top actions they want from government are clarity on how AI can and cannot be used, stronger enforcement of current laws, and active monitoring of AI systems. Older New Zealanders place especially high importance on clear accountability when AI systems cause harm or make mistakes. Data sovereignty appears to be a more secondary trust lever overall, but it matters more to those further along the adoption curve, indicating that local data processing and storage capabilities may become more important to New Zealanders in future.

Areas AI can Positively Impact NZ

#1 Boosting Productivity

Areas AI can deliver the greatest productivity gains

New Question Added in 2026**

Q. In which areas do you think AI could deliver the greatest productivity gains in New Zealand? (Select up to 3).**
Base: Those who selected 'Boosted productivity' as a positive impact 2026 (n=408)

INSIGHT 5 AI's value case is narrowing as expectations of guardrails rise

Can the Government shift New Zealanders' Trust in AI?

"If the government was to set clearer expectations and standards for how AI should be used, I would be more likely to trust AI - powered products and services."

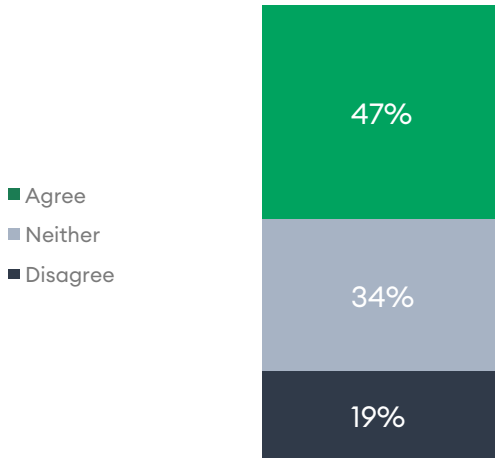

New Question Added in 2026**

Q. To what extent do you agree or disagree with the following statement? If the government was to set clearer expectations and standards for how AI should be used, I would be more likely to trust AI - powered products and services.**
Base: Total Sample 2026 (n=1,001)

Government actions that can increase trust in AI

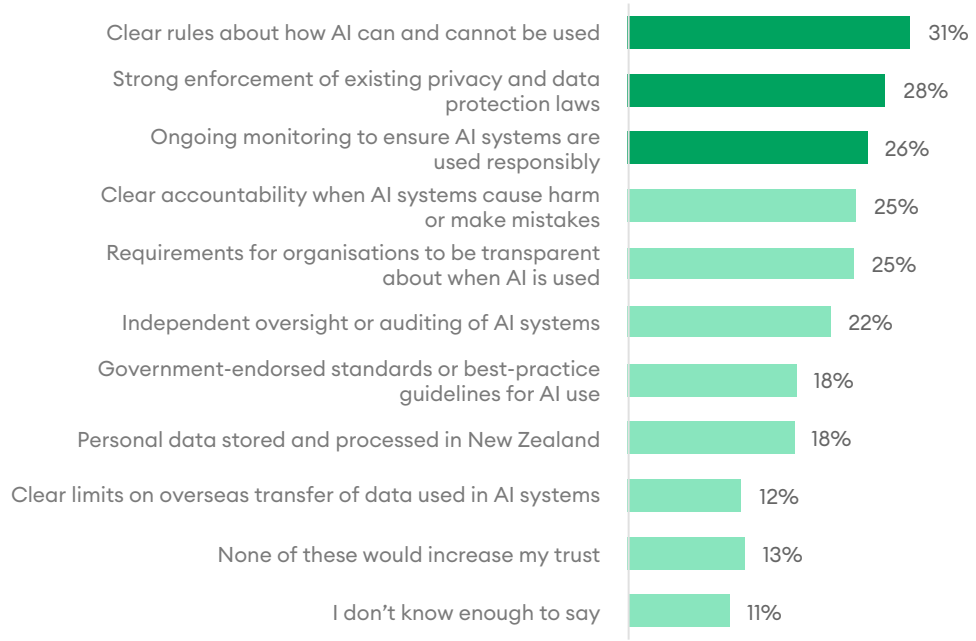

New Question Added in 2026**

Q. Which of the following government actions, if in place, would have the biggest impact on increasing your trust in AI - powered products and services? (Please select up to 3)**
Base: Total Sample 2026 (n=1,001)

INSIGHT 5 AI's value case is narrowing as expectations of guardrails rise

This is where the national productivity conversation gets more interesting – and more uncomfortable. As a country, New Zealand clearly wants AI to help solve efficiency and productivity problems. But New Zealanders do not appear willing to accept just any version of AI in exchange for cost savings. For many groups, especially women and younger Kiwis, protecting cultural values and identity ranks alongside, or above, economic growth and productivity. AI Actives are more likely to favour growth. Dabblers and Avoiders place more emphasis on wellbeing, fairness, or opt out of the value proposition entirely. That means the future of AI in New Zealand cannot be framed only as a race for efficiency. It also has to answer a cultural and social question: what kind of outcomes are we trying to create, and for whom?

Sustainability sharpens that challenge again. Forty-five percent of New Zealanders say they have concerns about the environmental impacts of business AI use. This rises to 63% among 18–24 year olds. Concerns are fairly consistent across the adoption curve, which means this is not just a fringe issue for sceptics. The dominant concern is the use of resources, particularly electricity. People say they would be more reassured if businesses took visible steps to reduce energy use or use renewable energy. But even here the signal is sobering: 30% of those with major sustainability concerns say that nothing would address them. This reinforces a broader pattern: expectations are rising faster than reassurance can keep up.

The strategic takeaway is that AI's value case is narrowing at exactly the moment the guardrail case is widening. That is not a reason to slow ambition. It is a reason to get more serious. In 2026, the organisations that lead will not just be the ones deploying AI. They will be the ones that can prove its value, explain its boundaries, show its oversight, and account for its broader costs.

Areas of Concern

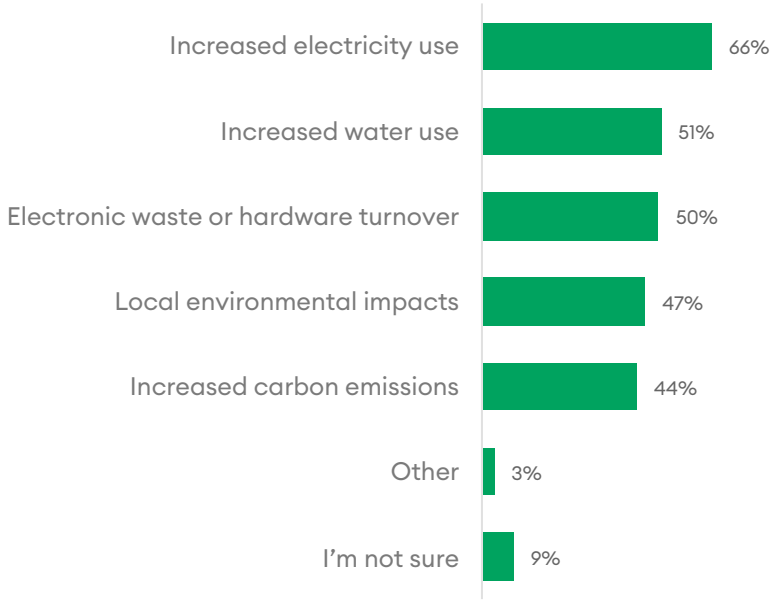

New Question Added in 2026**

Q. Which of the following are you most concerned about? (Select any that apply).**
Base: Those who are concerned about sustainability or environmental impacts (n=438)

What this means for Kiwi organisations and government

The macro challenge has changed. A year ago, the central question was how quickly AI would move into mainstream use. In 2026, the question is whether institutions can earn the confidence required for AI to go further.

New Zealand has a genuine opportunity here. AI could help address productivity constraints, improve service delivery and open up new forms of value. But that future will not be built on technical capability alone. It will be built on public legitimacy.

Leaders now need sharper answers to five questions:

1. **Where should AI assist, and where should it never decide alone?** This research makes clear that high-stakes settings demand visible human accountability, especially in healthcare, finance and government. If leaders cannot define those boundaries clearly, the public will assume they are not being managed.
2. **What proof of human oversight can a customer actually see?** Trust is not built by saying a human is in the loop somewhere. It is built when people can reach that human, understand what they are responsible for, and feel there is still a path to judgment when AI falls short.
3. **How will mistakes be detected, owned and put right?** If 70% of users are already experiencing issues, then incident management, escalation and accountability are now core parts of AI strategy. A bad deployment is not a side problem. It is a trust event.
4. **What public or customer value are you delivering beyond cost reduction?** New Zealanders still see the case for efficiency. But if AI is sold only as a tool for lower costs, it will struggle to earn broader legitimacy. The public wants productivity, but it also wants wellbeing, fairness, privacy and social benefit.
5. **Who is being left behind – and what is your plan for that?** Digital capability gaps, data sovereignty concerns and environmental impacts all sit inside the trust conversation now. If AI advances faster than public inclusion and safeguards, the country will widen the very divides it hopes technology can help close.

For government, the message is equally direct. The public is not asking for vague reassurance.

It is asking for clearer expectations, stronger protections, active monitoring, and visible accountability when harm occurs. That does not require government to solve every question before innovation can move.

But it does require setting credible rules of the road. Trust will not be rebuilt by optimism alone. It will be rebuilt by structure.

For any organisation helping shape New Zealand's connected future, the implication is not to claim perfection.

It is to lead with seriousness. Responsible AI is no longer just a governance topic at the edge of deployment. It is becoming a brand, customer and social licence issue at the centre of it.

Conclusion

New Zealand has moved beyond the first phase of its AI story. Awareness is high. Experimentation is real. But the public is no longer prepared to give AI the benefit of the doubt by default.

This research shows a country that is open to usefulness, but increasingly skeptical of hype. New Zealanders are interested in AI driving productivity, but more demanding about proof.

They are willing to engage, but only if human judgment, privacy, accountability and public safeguards remain intact.

That should be taken seriously, not defensively. It means New Zealand still has room to shape its AI future before trust erodes further. The opportunity is not gone. But it is becoming more conditional. The organisations that succeed in the next phase will not be those that treat trust as a communications layer added after deployment. They will be those that build it into design, governance, operations and customer experience from the start.

There is also a broader structural context beginning to emerge behind these trust questions. Much of the advanced AI capability New Zealand relies on today is developed and operated offshore, often by a small number of global platforms. As AI becomes more embedded in economic activity and public services, this concentration raises new considerations for smaller economies – including questions of dependency, resilience and where local capability matters most.

This does not suggest New Zealand needs to build everything itself. But it does point to a more deliberate set of choices ahead: where to adopt global AI capabilities, and where greater control, transparency or local oversight may be required to maintain trust and confidence over time.

The national question, then, is not whether New Zealand wants AI.

It is what kind of AI future New Zealand is willing to trust.

Methodology and role of AI

The 2026 One NZ AI Trust Report was conducted by Perceptive using an online survey of New Zealanders aged 18 and over.

The research used a nationwide sampling framework and results were weighted to Statistics New Zealand census data, including gender, age and location, to produce a nationally representative sample. A total of 1,001 responses were collected.

Fieldwork took place between 6 March and 20 March 2026. The study has a margin of error of plus or minus 3%.

PERCEPTIVE
PART OF CLEMENGER UNLTD

This report was developed through a human-led, AI-assisted process using a range of tools, with ChatGPT being used most extensively.

AI was used to support drafting, synthesis and editing, while all analysis, review and final approval remained the responsibility of the One NZ team.

The use of AI throughout the report development process was undertaken in line with One NZ's Responsible AI standards.

one.nz
LET'S GET CONNECTED

Ngā mihi

For more information, please visit our website:

one.nz/business