

Appendix B

Out of scope

From: LensenMcGavin AI S 9(2)(a)
Sent: Wednesday, 6 May 2026 12:32 pm
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Subject: Re: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance

Kia ora Alex,

No notes from our end. Absolutely brilliant.

Without wanting to sound too much like a sycophantic chatbot myself, you and your team should be very proud. The improvement from the first draft you sent to this version is remarkable. If you update the Te Tiriti considerations and the risk-based assessment tooling sections, you will have produced the gold standard for public sector AI guidance.

Out of scope

Out of scope

Out of scope

Out of scope

Warm regards

S 9(2)(a)

On Wed, 6 May 2026 at 08:33, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a)

Please find attached our final draft of the AI guidance and a draft mock-up of a 4 page summary to accompany the guidance (my comms team still needs to review this also). If you're able to provide a final quick review for technical accuracy, within the contracted amount already paid, that would be much appreciated.

As you'll notice most of your feedback has been incorporated into the final draft. Where feedback hasn't been directly added, such as CBA and risk-based assessment tool, this is because further work is happening in this space and will likely be incorporated into updates of the guidance.

To help with your quick review, the below content has been specifically added:

- Acknowledgement of LensenMcGavin AI on page 2
- Emphasis on AI literacy and corrected interpretation of the Public Service Census stats throughout
- A simple explanation of the two types of AI added on pages 4 and 5
- Automation bias on pages 9, 21 and 23
- Te Tiriti o Waitangi considerations on pages 28 and 29. This will be expanded following sector feedback and support from Te Puna Kokiri in future updates.

Out of scope




Kia pai tō rā,

Alex

Alex McMinn

Head of Regulatory System Capability

Ministry for Regulation | Te Manatū Waeture

īmēra: alex.mcminn@regulation.govt.nz | **waea pūkoro:** 

Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: LensenMcGavin AI S 9(2)(a)
Sent: Monday, 13 April 2026 12:31 pm
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Subject: Re: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance

You don't often get email from S 9(2)(a) [Learn why this is important](#)

Kia ora Alex,

Happy to help! Just let us know when and if.

Warm regards

S 9(2)(a)

LensenMcGavin AI

On Mon, 13 Apr 2026 at 08:04, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a),

Thank you for the fulsome review and detailing what works well, what needs expanding/re-writing and where the gaps or errors are. I'll be making my way through the updates this week as we plan to publish late April, following discussions with Minister Seymour's office.

If possible, I may come back to you with just a further review of the paragraph where I build on the explanation of the different AI models for technical accuracy.

Your invoice has been sent on for payment.

Kia pai tō rā,

Alex

Alex McMinn

Head of Regulatory System Capability

īmēra: alex.mcminn@regulation.govt.nz | waea pūkoro: S 9(2)(a)

Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: LensenMcGavin AI S 9(2)(a)
Sent: Wednesday, 8 April 2026 2:32 pm
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: S 9(2)(a)
Subject: Re: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance

You don't often get email from S 9(2)(a) [Learn why this is important](#)

Kia ora Alex,

Thank you again for engaging LensenMcGavin AI in the external review of your draft AI guidance for regulators. We have now completed our review and attach our report.

As the report indicates, we support much of the guidance but also highlight a few key areas for improvement. We remain available to assist further in this workstream as needed.

Best,

S 9(2)(a)

From: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Date: Mon, 23 Mar 2026 at 11:46
Subject: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance
To: S 9(2)(a)

Thanks S 9(2)(a). Yes, GST will be added to the total amount.

Thanks for signing the attached contract. I've now attached the draft AI guidance for your full review.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: S 9(2)(a)
Sent: Monday, 23 March 2026 9:30 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Subject: Re: Engagement – Peer Review of AI in Regulatory Practice Guidance

Kia ora Alex,

Looks good on our end. I presume the GST will be on top?

Best

S 9(2)(a)

On Fri, 20 Mar 2026 at 16:54, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a)

Can you please review the attached lite model contract and send back to me and then once I have checked by my legal team, I can then formally send the contract and guidance document to you by Monday/Tuesday next week.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



**Ministry for
Regulation**



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: Alex McMinn

Sent: Thursday, 19 March 2026 4:28 pm

To: S 9(2)(a)

Cc: S 9(2)(a)

Subject: RE: Engagement – Peer Review of AI in Regulatory Practice Guidance

Kia ora S 9(2)(a)

I should have the contract and draft ready for you next week. Thanks for checking in.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



**Ministry for
Regulation**



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: S 9(2)(a)
Sent: Wednesday, 18 March 2026 11:08 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: S 9(2)(a)
Subject: Re: Engagement – Peer Review of AI in Regulatory Practice Guidance

Hi Alex,

Hope you're well. Just checking back in on this.

Best

S 9(2)(a)

On Tue, 3 Mar 2026 at 08:16, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a),

That's great to hear — thank you. I'm currently revising the draft guidance following external feedback, including from colleagues at GCDO.

All going to plan, I expect to send you the CMC with a final draft in a couple of weeks. As you may have gathered, I'm also flexible on the publishing date to ensure feedback is received and considered meaningfully.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



**Ministry for
Regulation**



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: S 9(2)(a)
Sent: Wednesday, 25 February 2026 9:54 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: S 9(2)(a)
Subject: Re: Engagement – Peer Review of AI in Regulatory Practice Guidance

Mōrena Alex,

Thanks for this, that is great to hear.

As the nature of the work is relatively uncomplex, we would be happy to use either your standard consulting agreement or the GMC-lite. A similarly 'lite' proposal document which confirmed the scope and agreed fee of S 9(2)(b)(ii) would be fantastic.

Given the publication time frame, we would be happy to commence as early as March 2nd and can prioritise delivery to ensure completion no later than March 16th. But, we can be reasonably flexible to suit your requirements.

Look forward to working with you as well. Let me know if there's anything more specific you need from us at this stage.

Ngā mihi nui

S 9(2)
(a)

On Tue, 24 Feb 2026 at 14:58, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a),

Just a quick note to let you know I now have approval to proceed with engaging LensenMcGavin AI to undertake the external peer review of our *AI in Regulatory Practice* guidance.

The scope is advisory only and focused on:

- clarity and coherence of the draft
- framing for senior regulatory leaders
- regulatory defensibility, including judgement, governance, and accountability settings
- identifying any material gaps or areas of overreach

We're working to a late March publication timeframe, so would be looking to share the near-final draft with you shortly and receive written feedback within an agreed window. Could you let me know what you need from us to formalise the engagement (e.g. proposal, confirmation of scope/fee, contract details), and your availability over the next couple of weeks?

Appreciate your support on this and look forward to working together.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



 *Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)*

From: S 9(2)(a)
Sent: Tuesday, 24 February 2026 10:21 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Subject: Re: 20260128 AI Draft Guidance - version for external feedback 0.2

Hi Alex,

Thanks for the update. Will keep my eyes peeled.

Have a lovely week!

Best

S 9(2)(a)

On Tue, 24 Feb 2026 at 09:14, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a),

I'm just working through a possible procurement exemption process. I hope to have an update later this week.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



 Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: S 9(2)(a)
Sent: Monday, 23 February 2026 9:29 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Subject: Re: 20260128 AI Draft Guidance - version for external feedback 0.2

Morning Alex,

Hope you're well. Just wondering where we got to on this?

Best

S 9(2)(a)

On Thu, 12 Feb 2026 at 13:04, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Thanks S 9(2)(a), that's great to know. Timeframe and price sounds good.

As it's a one off specialised piece of work I can likely get an AoG procurement exemption.

I'll talk to my procurement lead and get back to you on this.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | **M:** S 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

From: S 9(2)(a)
Sent: Thursday, February 12, 2026 12:15 PM
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: S 9(2)(a)
Subject: Re: 20260128 AI Draft Guidance - version for external feedback 0.2

Hi Alex,

Good to know. There are a number of passages and turns of phrase throughout the document that come across as very 'ChatGPT'.

Yes, we would be happy to provide peer review service. We are not currently registered on the AoG procurement panel. However, we could register if that made things smoother on your end. I'm not sure on timelines for this.

For this type of work we generally charge a flat fee depending on scale and complexity. In this instance it would be S 9(2)(b)(ii). This would comprise 20 hours of work inclusive of edits, feedback and review. We would be positioned to have the work delivered by the end of February.

Happy to discuss further.

Best

S 9(2)(a)

On Thu, 12 Feb 2026 at 11:15, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Thanks § 9(2)(a)

AI (ChatGPT specifically) was used mainly as an editing assistant. Collating our interview notes into themes and ideas, plain language and summarising.

As part of planning for external peer review, is this something LensenMcGavin AI does? If so, could you give me more information on this, such as registered on the AoG Procurement Panel, timeframes and cost? If not, is there someone else you could suggest for this.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: § 9(2)(a) | **M:** § 9(2)(a) | **E:** alex.mcminn@regulation.govt.nz



 *Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)*

From: § 9(2)(a)
Sent: Thursday, 12 February 2026 10:06 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: § 9(2)(a)
Subject: Re: 20260128 AI Draft Guidance - version for external feedback 0.2

Hi Alex,

Thanks for this. It was an interesting read, to say the least. I'm very curious to know the extent to which AI was used in the drafting process.

Best

S 9(2)(a)

On Wed, 11 Feb 2026 at 16:32, Alex McMinn <Alex.McMinn@regulation.govt.nz> wrote:

Kia ora S 9(2)(a),

I am writing to give you a brief update on the development of the Ministry for Regulation's AI guidance.

I want to thank you all for your time, insights, and thoughtful contributions you provided throughout the research process. Your perspectives have been instrumental in shaping the direction, scope, and practicality of the guidance. The feedback you shared, both through discussions and written input has helped ensure the work reflects real needs, diverse expertise, and the realities of how AI can be applied in regulatory contexts.

As part of the development process, we have deliberately aimed to adopt a style that is clear, practical, and accessible for managers, recognising the need for guidance that supports decision-makers.

Over the next two weeks, I will be finalising the attached draft of the guidance, including having it externally peer-reviewed. While the main feedback phase is complete, I am sharing this draft as an FYI. If you notice anything that you think would benefit from clarification, or correction, please feel free to let me know.

I will keep you updated on when the guidance will be publicly available.

Kia pai tō rā

Alex McMinn

Head of Regulatory System Capability

Reviews and System Capability

Ministry for Regulation

P: S 9(2)(a) | M: S 9(2)(a) | E: alex.mcminn@regulation.govt.nz



**Ministry for
Regulation**



Like to book a meeting with me? Meet [in-person](#) or [Teams](#) / [Quick call](#)

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.


.....
Confidentiality notice: This email may be confidential or legally privileged. If you have received it by mistake, please tell the sender immediately by reply, remove this email and the reply from your system, and don't act on it in any other way. Ngā mihi.

From: [Alex McMinn](#)
To: reginvoices@cass.govt.nz
Subject: FW: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance
Date: Monday, 13 April 2026 8:06:00 am
Attachments: [LensenMcGavin AI Invoice #107.pdf](#)


Out of scope



From: LensenMcGavin AI 
Sent: Saturday, 11 April 2026 8:00 am
To: Alex McMinn <Alex.McMinn@regulation.govt.nz>
Cc: 
Subject: Re: FW: Engagement – Peer Review of AI in Regulatory Practice Guidance

You don't often get email from . [Learn why this is important](#)

Out of scope




On Wed, 8 Apr 2026 at 14:31, LensenMcGavin AI  wrote:

Kia ora Alex,

Thank you again for engaging LensenMcGavin AI in the external review of your draft AI guidance for regulators. We have now completed our review and attach our report.

As the report indicates, we support much of the guidance but also highlight a few key areas for improvement. We remain available to assist further in this workstream as needed.

Best,




LensesMcGavin AI Independent Peer Review: Draft AI Guidance for Regulatory Practice

Prepared by: s 9(2)(a) and s 9(2)(a) (Directors, LensesMcGavin AI)

Prepared for: Ministry for Regulation (MfR)

Date: April 2026

Document Status: First Draft Review

1. Introduction and Context

The Ministry for Regulation (MfR) has developed draft guidance intended to support senior leaders and management teams in regulatory organisations as they navigate the integration of Artificial Intelligence (AI) into regulatory practice. LensesMcGavin AI was engaged to conduct an independent peer review of this draft document. The purpose of this review is to provide objective feedback on the clarity, framing, and regulatory defensibility of the guidance, and to suggest improvements that strengthen its practical applicability for regulatory leaders.

We are available to review subsequent drafts or discuss these recommendations in detail with the drafting team.

2. Scope and Methodology

Our review focused on evaluating the draft guidance against current best practices in AI governance, public sector digital adoption, and regulatory risk management. The assessment prioritised the following criteria, as indicated by MfR:

1. **Regulatory Defensibility:** Ensuring the guidance adequately addresses the unique legal and ethical obligations of regulators, particularly regarding transparency and accountability.
2. **Practical Applicability:** Assessing whether the advice provides actionable, realistic steps for public sector leaders.
3. **Technical Accuracy:** Evaluating the document's presentation of AI capabilities and limitations.

This review is based on the draft provided to us in March 2026. It is designed to offer strategic recommendations for the document's refinement and does not constitute formal legal advice.



3. Executive Summary

The draft guidance provides a timely and highly relevant foundation for regulatory agencies exploring AI. It identifies that integrating AI goes far beyond standard software or IT deployment, instead viewing this process through a change management lens. The document aims to centre human decision-making through Human-in-the-Loop (HITL) processes, and takes a principled stance on accountability, auditability, review, establishing ethical boundaries, and integrating Māori data sovereignty principles.

The document forms the basis of what has the potential to be a particularly strong, pragmatic, and actionable piece of guidance, but would benefit from improvements in a few key areas. Most notably, the guidance frames AI as being highly adopted but not yet transformational. This is at odds with recent Public Sector Commission reporting that shows AI is only utilised by 12% of officials on a daily or weekly basis.

To strengthen the guidance, we suggest adjusting the underlying tone by placing a greater priority on foundational AI literacy as an essential prerequisite for safe operational use. Further improvements would include clarifying technical definitions to differentiate between predictive and generative AI, explicitly addressing the established risk of automation bias within HITL processes, reinforcing statutory accountability by linking transparency to the Official Information Act, and enhancing procurement advice to ensure due diligence on data provenance and contractual bias levels. The crucial point that accountability remains firmly with the organisation and legally authorised decision-makers should also be made much more prominent.

The inclusion of Te Tiriti o Waitangi, Māori data, and Māori data sovereignty concepts is appropriate, but is currently very high-level in nature. These concepts should be considered in more detail, drawing on the assistance of Māori experts and stakeholders. The current lack of depth risks the guidance being seen as passively referencing these issues rather than providing meaningful insights that regulatory leaders can consider implementing.

4. Key Thematic Findings

Our analysis of the document identified several broad thematic areas in which improvements could be made. These areas are detailed in the following subsections.

4.1 Strategic Pacing and AI Literacy

The current draft conveys a significant sense of urgency regarding the adoption of AI. It appears to mischaracterise New Zealand's current position as one of high AI adoption and low transformation. Conversely, available literature, including sources cited within the document, suggests that the country is currently experiencing a period of low adoption and low transformation. We therefore recommend placing a greater priority on the development of AI



literacy. This would include developing officials' understanding of AI's technical dimensions, specifically how training data impacts outputs and the mechanisms by which models generate responses. Advising regulatory leadership to dedicate time to establishing this foundational understanding prior to pursuing widespread operational deployment will produce AI systems that are safer, of higher quality, and more sustainable for long-term use.

4.2 Technical Grounding

The guidance uses the term "AI" broadly, often appearing to refer specifically to Large Language Models (LLMs). The guidance would be strengthened by clarifying the specific types of AI being discussed, particularly differentiating between predictive AI (systems that classify, score, or forecast outcomes based on existing data) and generative AI (systems that create new content, such as text, code, or images). These different categories carry distinct regulatory requirements and risk considerations. For example, the use of predictive AI requires scrutiny of bias in training data and fairness of outcomes, while generative AI systems need to be carefully managed to maximise factual accuracy (reducing hallucinations) and to respect intellectual property rights.

The guidance indicates it will deliver a simple plain English explanation of what AI is (and isn't), but our review could not find this explanation. There is some information on how AI works under the 'Understanding AI in a regulatory context' heading, but it is lacking in detail. There is also a lack of depth in how the guidance discusses the technical limitations and risks of AI. There is currently no definition of phenomena such as "hallucinations", and limited explanation of risks and how they materialise. Including an explicit description of these important concepts would assist regulators in understanding in a more practical manner how these systems work and what their limitations are.

4.3 Risk Management

The guidance provides a checklist which is intended to be used by regulatory leaders when exploring the possibility of using AI in regulatory work. This is a useful and practical tool that could be rolled out as part of a standardised risk management process for regulatory leaders.

There are very good processes in this checklist that are not well represented in the body of the document, such as the suggestion to undertake a risk-based assessment prior to adoption. A cost-benefit analysis would also be another useful addition to this checklist, which would allow regulatory leaders to determine whether a specific problem truly necessitates an AI solution, or if traditional automation would be safer and more effective.

It may also be helpful for this checklist to be supported by a risk management schema to assist regulatory leaders when proposing to use AI in higher-risk settings.



4.4 Human-in-the-Loop & Automation Bias

While the guidance correctly positions human-in-the-loop (HITL) processes as a safeguard against AI risk, it does not explore the established risk of automation bias, where human reviewers routinely defer to machine outputs without considering their own judgement. This phenomenon can undermine the effectiveness of HITL. We suggest explicitly addressing this risk and advising agencies to design systems that actively encourage reviewers to challenge AI outputs and to implement oversight mechanisms to challenge and review the human reviewers as well. Otherwise, it is possible that regulators may fall into the trap of overreliance on HITL.

4.5 Transparency

The guidance notes the importance of transparency but misses an opportunity to connect this to existing statutory obligations. Referencing the Official Information Act, specifically the public's right to an explanation regarding decisions that affect them, would reinforce the necessity of explainable AI in regulatory settings. Similarly, the guidance correctly identifies the importance of review, appeal and audit trails. However, it does not offer any specifics as to what these processes might look like. In particular, a standardised process for auditing across regulators would reduce risk and improve service delivery considerably.

4.6 Te Tiriti o Waitangi, Māori Data, and Māori Data Sovereignty

The guidance contains a relatively small number of references to Māori data, Māori data sovereignty and Te Tiriti o Waitangi. We commend the inclusion of these vitally important concepts. However, in our view the extent to which these concepts are explored and explained in the document is insufficient. There is a real risk that in their current state, the concepts will function merely as lip-service and sign posting, as opposed to actionable and practical additions that would markedly improve outcome for Māori. We view this as particularly important, as many of the government outputs which form the AI policy/guidance ecosystem neglect (and have historically neglected) to properly integrate these principles despite the fact that Māori will disproportionately feel the negative impacts of AI. In our view, this guidance could be an opportunity to meaningfully progress and explore Te Tiriti, Māori data and Māori data sovereignty as they relate to AI and regulatory practice. As such, we strongly recommend that MfR engages with experts and leaders in this domain to properly inform these components of their guidance.



5. Detailed Review and Suggested Amendments

The table below outlines specific observations from the draft text and provides corresponding recommendations for refinement.

Section & Page	Current State / Issue	Suggested Amendment /Recommendation/Comment
Who is this guidance for?", page 4.	The document talks about how it complements existing government-wide frameworks, but these are not linked until later. Mixture of terms "useful links" and "references".	Tidy up relationship between MfR AI Guidance and existing AI ecosystem. Given the current state of AI policy environment, it is likely that a regulator will have to consult at least some of these documents.
Who is this guidance for?", page 4.	"This creates distinct considerations for how AI is used in regulatory settings."	Reworking this sentence to include a mention of 1 or 2 prominent AI use cases in regulatory settings would be of benefit.
Why AI matters for regulation?", page 4.	Bullet list mentions a number of areas where AI can support regulatory practice	This list could be expanded or contextualised with information on how these actions can support regulatory practice.
"Understanding AI in a regulatory context", page 5	At the outset the document promises a simple plain English explanation of what AI is and isn't. In its current form this definition is insufficient. "Before diving into how to use AI..."	Introduce a more expansive explanation of how AI works to better inform users on where and how it may be used (i.e. AI literacy)
"Understanding AI in a regulatory context", page 5	As above.	The current explanation takes 'AI' to mean LLMs. Is this guidance for LLMs only? Imprecise language is an issue here.
"Understanding AI in a regulatory context", page 5	"That is why regulators need strong Human-in-the-Loop (HITL) arrangements."	The guidance should also highlight that a key part of creating HITL arrangements is creating review and audit processes to reduce the likelihood of human error and impacts of automation bias, etc.
"Understanding AI in a regulatory context", page 5	"Building clear human oversight and decision-making control into AI-supported processes means AI support"	Should read: "Building clear human oversight and decision-making control into AI-supported processes means AI supports"
"Understanding AI in a regulatory context", page 5	"The following resources	This could be expanded further to show



regulatory context”, page 5	provide useful information about Māori data and Māori data sovereignty”	how these principles may apply to using AI in a regulatory context. These considerations could also be mentioned elsewhere in the document to provide context. Best practice would suggest external consultation, otherwise the section risks merely being “lip-service”.
Section one: set the foundations, page 6.	“It is more likely to expose and amplify existing problems”	It would be good to expand on these problems. “Such as...”.
Section one: set the foundations, page 6.	For example, if inspection, licensing or compliance datasets are incomplete, inconsistent, or poorly structured AI systems may reinforce existing weaknesses...”	It would be good to expand on what these weaknesses might be.
Section one: set the foundations, page 6.	“AI also changes the nature of regulatory risk....Despite the change in risk accountability does not move to the tool or vendor. It remains with the organisation and the legally authorised decision-makers.”	The fact that organisations and legally authorised decision-makers retain accountability is a very important point and should be much more prominent.
Section one: set the foundations, page 6.	The document frames AI literacy primarily around knowing how to use tools, with an underlying urgency to adopt quickly.	Shift the emphasis to highlight that effective AI literacy requires an understanding of technical mechanisms. Advise regulators to prioritise this foundational literacy before encouraging broader implementation beyond prototypes/sandboxing.
Section one: set the foundations, page 6.	“This is why AI adoption should not be treated as an IT project. It is a change programme...”	The messaging here is good (contrasting AI integration with a software upgrade). Framing this in terms of organisational buy-in is more appropriate than change management.
Start strong: strategy, governance and guardrails, page 7.	“AI adoption is most successful when it is treated as a change management with a strategic regulatory	Missing word ‘programme’ after change management. Double space between IT and project. Unnecessary repetition of above change management content. Instead, this could be an appropriate



	capability focus, not a stand-alone IT project.”.	location to discuss the importance of understanding the limitations of the technology, especially in its application to regulatory environments.
Start by setting three foundations, page 7.	1. Align AI to purpose and priorities	In this section, there should be the addition of a cost-benefit analysis or a risk-based decision framework. For instance, what is the expected RoI on AI use, what is the risk profile of using AI vs BAU.
Start by setting three foundations, page 7.	“This keeps adoption grounded in outcomes, rather than driven by novelty or vendor momentum.”.	Vendor momentum could be changed to less emotive language, “vendor promises” comes to mind.
Put governance and decision rights in place early, page 7.	“Every AI assisted regulatory activity should have clear ownership, decision rights, and escalation pathways.”.	Very good.
Put governance and decision rights in place early, page 7.	“Including an appeal to a human where AI has influenced outcomes.”.	This should be very clearly sign posted, and displayed in any communications about that outcome. This process should also be routinely audited to ensure it is working, and that individuals feel empowered to and are easily able to challenge AI influenced decisions or outcomes.
Designing for learning and adaptability, page 7.	“Regulators should take an iterative approach..”.	An ‘iterative approach’ should start small and emphasise getting regulators comfortable with AI. Once they understand how the technology works, they can then move to begin looking at how it might work within their existing workflows.
Designing for learning and adaptability, page 7.	“Many regulators benefit from establishing an AI governance group”.	This should be a much stronger recommendation. Each regulator (or groups of regulators) should have some form of AI governance resource. Does not necessarily have to be an in-house function, it could be outsourced.
Designing for learning and adaptability, page 8.	Finally, build in regular review points to check whether systems remain fit for purpose...”.	Very good. A key part of this process could also be the review of the HITL, including the humans who are in that loop. Very important to solve for / design with human error in mind. HITL is not a silver bullet.
Practical steps for	Apply relevant government	Is this the guidance linked at the start of



engaging AI vendors, page 8.	guidance.	the document? There is a question as to the suitability of some of the public sector AI guidance in serious procurement discussions.
Do vendor due diligence, page 8.	All of subhead 3.	We think more information in this section would be useful. For instance, it could be good to gain an understanding of the provenance of the data used in the model depending on use-case. Further, testing and benchmarking data and contractual requirements as to levels of bias could be good ways to improve interpretability. It may be pragmatic to emphasise outputs rather than models themselves. Procurement teams should endeavour to get as much information as they can, though it is unlikely they will be granted access to weights and other internals. The use of some manner of risk based assessment is likely best practice, higher risk applications should require models to have a much higher standard of interpretability.
Do vendor due diligence, page 8.	"Including sovereignty where relevant".	Sovereignty should always be a relevant consideration.
Do vendor due diligence, page 8.	"How model is tested, updated, and monitored".	Consider seeking disclosure as to training and testing datasets, or getting some manner of explainer at a minimum. There is nothing in this guidance about model training processes.
Identifying and manage key risks early, page 9.	"Common risks include bias, security weaknesses, and vendor lock-in.".	This section needs much more information, and more explicit guidance as to risks from models.
Lock in accountability through the contract, page 9.	"Include clauses on..".	Procurement teams may want to include clauses which set costs for models. If frontier model prices increase there will be potentially drastic price increases down stream, so vendor selection is paramount in this regard.



Lock in accountability through the contract, page 9.	Data portability and access to outputs generated using agency data.	It is important to get certainty as to where data is stored.
AI is a change process: getting people ready for AI, p9.	"AI is a change process, getting people ready for AI."	We like this framing as it recognises that it goes beyond mere tech deployment. However, calling it a change process does imply some amount of pre-determined outcome. We wonder if there's scope to suggest a consultation process (as part of good "change management") to find sticking points early on, or potential flags to reconsider the use of AI.
Key reasons why change management matters for AI adoption, page 9.	1. Cultural shift	A key aspect of this change management framing is AI literacy. Staff and stakeholders should be empowered to understand how and why AI does the things it does. This understanding will support HITL and critical thinking. As above, you do not want to be over-reliant on HITL, you need to support it with processes which reduce human error. Current research indicates that there is a considerable amount of groundwork required to bring staff along with you, and that it is very hard to get RoI from AI without this.
2 Role definition, page 9.	"Over time, some staff will spend less time on manual tasks and more time on oversight, assessment and interpretation."	This wording presumes that AI will be sufficiently good to replace a substantial amount of human work. Could tweak wording here.
3 Capability building, page 10.	"Regular use of AI for work remains relatively low."	Should add following stats from PSC report: 12% use AI daily or weekly. 67% have never tried using AI for work. Again, should emphasise the need for AI literacy as not just a standard software rollout / digital skill (sociotechnical element is very important.)
Practical steps (what regulatory leaders	"Practical steps (what regulatory leaders can do	Missing a word: "can do to ensure..."



can do ensure a successful AI rollout), page 10.	ensure a successful AI rollout)".	These aren't really practical steps, so much as they are a sign post to what practical steps could be. We would suggest bulking this section out and providing more explicit advice.
High adoption, low transformation, page 11.	"Pilot and iterate: start small with low risk use cases and scale only when value and safeguards are proven."	Very good.
Create feedback loops: listen to staff and affected groups and adjust quickly when risks or issues appear, page 10.	"Create feedback loops: listen to staff and affected groups and adjust quickly when risks or issues appear".	Very good.
2.1 High adoption, low transformation, page 11.	"This creates an adoption paradox: use is high..".	The above cited PSC report states that only 12% of people use AI daily or weekly,
High adoption, low transformation, page 11.	"Without deliberate leadership and workflow redesign, AI remains helpful at the margins rather than allows us to achieve regulatory outcomes more efficiently and effectively."	Word missing: "AI remains helpful at the margins rather than allowing..." Research suggests that a large degree and organisational and cultural change is required to see benefit from generative AI systems, a large number (up to 95% of organisations are getting zero RoI)
1. Treat early adoption as the starting point, not the destination, page 11.	1. Treat early adoption as the starting point, not the destination.	What is "the destination" here?
1. Treat early adoption as the starting point, not the destination, page 11.	"Widespread use of tools like Copilot builds familiarity."	Familiarity with the tool is only one facet. This needs to be combined with some understanding of technical aspects.
2. Create an authorising environment for safe experimentation, page 11.	"Leaders don't need certainty about how AI will be used."	We would suggest that this is rewritten to state "there should be freedom to experiment safely". Leaders do probably need some certainty about how AI will be used.
3. Look for regulatory use cases, page 12.	"These questions point to safe starting points for AI support."	A key part of this inquiry should be whether or not solving the task actually requires AI, or if it could be supported by automation or



		some other such tool. These questions arguably do not point to “safe” starting points.
4. Start small and let progress compound, page 12	“Small, low-risk steps build capability and momentum, without needing major restructures or large technology programmes. Recognising the “high adoption, low transformation” phase helps regulators be honest about where they are, and what needs to change to unlock real value”.	Literature suggests that a large degree of organisational impetus, training and communication is required to make AI work. As supported abovementioned PSC study, high adoption is not here yet. It would be better to frame this as: a way to encourage/enable adoption is to start with small, low-risk steps.
5. Build confidence before chasing transformation, page 12.	“The early goal isn’t large-scale change”.	Does this imply that the end goal is large-scale change?
6. Learn from AI adoption across the regulatory system, page 12.	“Peer networks and communities of practice help build collective capability, reduce duplication, and develop a shared understanding of good practice.”.	This should also feed into revised guidance and frameworks. Lessons learned need to be formally shared.
2.2 Examples of AI in regulatory practice, page 12-14	Case Studies: AI in permissions & concessions, AI Volution - decision support.	We are very pleased to see some use cases detailed here. It would be highly beneficial to the Guidance if both a negative and high risk example were included to highlight risks and risk management strategies.
2.3 AI Literacy: safe and purposeful use, page 15.	“AI literacy turns early experimentation into confident, responsible use.”.	Very good to highlight the importance of AI literacy. Unsure the line between AI literacy and confident responsible use is this linear.
2.3 AI Literacy: safe and purposeful use, page 15.	“AI literacy does not require technical expertise.”.	We disagree here, we think that good AI literacy should develop some level of technical expertise insofar as how models work (abstractly), why they produce the outputs they do, and their limitations.
2.3 AI Literacy: safe and purposeful use,	“What AI can and cannot do”.	Arguably, this requires some technical knowledge.



page 15.		
2.3 AI Literacy: safe and purposeful use, page 15.	"Stalled progress, leaving regulators stuck in "high adoption, low transformation"".	Unclear. The PSC report indicates that there is not high adoption. Does adoption in this case mean that regulators / agencies feel pressured to adopt?
2.3 AI Literacy: safe and purposeful use, page 15.	"AI literacy is not the end goal. It is the starting condition for using AI well, keeping humans in control, and integrating AI safely into regulatory practice".	This is not very strong writing. AI literacy should be a constant.
2.4 AI as a support for regulatory thinking, page 15.	Whole section	As a general statement, the writing in this section is currently weaker than in the rest of the document.
2.4 AI as a support for regulatory thinking, page 15.	"Research in cognitive science".	It would be beneficial to provide a link to a supporting study. Such as, Your Brain on ChatGPT .
2.4 AI as a support for regulatory thinking, page 15.	"The risk is subtle but real" the tools improve, while the people responsible for decisions become less confident and less capable.".	Consider rewriting.
2.4 AI as a support for regulatory thinking, page 16.	"When AI outputs are accepted without challenge, three risks arise.".	It would be good to link these risks to an eroding of public trust.
2. Illusion of understanding. Page 16.	"AI outputs often sound confident, even when they are incomplete or wrong. Without strong domain knowledge fluency can be mistaken for accuracy.".	It would be good to include a summary or definition of automation bias here.
What good looks like, page 16.	"Regulatory capability grows when AI sharpens judgement, not when it substitutes for it.".	Rewrite for clarity.
Leadership implications, page 16.	"Used carefully, AI can strengthen regulatory practice over time. Used poorly, it can de-skill the workforce and undermine the legitimacy of decisions.".	Repetition.



Keep humans central, page 16.	“This is where concepts like HITL, human oversight, and the right to appeal to a human become essential.”.	Section 23 of the Official Information Act 1982 is applicable here.
2.5 Human-in-the-Loop: keeping decision-making human, page 17.	“AI can help by scanning large volumes of information, identifying patterns, or highlighting risks. But it cannot understand context, apply discretion, or weigh legal and ethical considerations. HITL ensures that practitioners understand, question, and, where necessary, challenge or override AI outputs before a decision is made.”.	This section needs to be bulked up. The guidance is severely lacking a serviceable explanation of what AI does and how it does it. It seems contradictory that this model of HITL ensures practitioners understand model outputs, but the guidance has very little requirements for technical understanding. The HITL described below the figure, appears to be more onerous than as described in the figure. Again, Section 23 of the Official Information Act 1982 is applicable here and should probably be built into the process.
2.5 Human-in-the-Loop: keeping decision-making human, page 18.	“Human judgement therefore influences AI both in how systems are designed and how outputs are used in practice.”.	Reinforces the need for the guidance to advocate for AI literacy training.
2.5 Human-in-the-Loop: keeping decision-making human, page 18.	“AI adds value only when regulators can explain how its insights were used, why they were accepted or rejected, and how the final decision was reached. AI may support decisions, but accountability remains firmly with people.”.	This is a legal obligation and regulators should be very clear about this.
2.6 Iterative unleashing: learning safely by doing, page 19.	2.6 iterative unleashing: learning safely by doing	The choice of this approach is not clearly justified.



2.6 Iterative unleashing: learning safely by doing, page 19.	“The goal is not to deploy AI quickly, but to build organisational judgement about where AI adds value, where it creates risk, and how it can be safely used over time.”.	This is good, but delivers different messaging to the language used earlier in the guidance (which was more pro-AI use).
How iterative unleashing works in practice, page 20.	“Examples include scanning public sentiment, organising large volumes of public information, or identifying themes in submissions or complaints. These are tasks where outputs can be checked easily, errors are unlikely to cause harm, and the work would otherwise consume significant staff time.”.	The public generally expects regulators to review submissions and complaints (i.e. with human eyes). The social license needs to be considered here, especially in higher risk regulatory settings.
How iterative unleashing works in practice, page 20.	“The aim is learning, not perfection. Sometimes the output is useful; sometimes it is not; sometimes it reveals underlying data, or process issues the organisation had not seen before. All of this is valuable insight.”.	Rewrite.
How iterative unleashing works in practice, page 20.	“Reversibility is critical. Each pilot is designed so it can be paused, rolled back, or stopped altogether. Ending a pilot is not a failure. It is part of building informed capability.”.	This section should be tied to sandboxing.
Linking back to HITL, page 20.	“AI supports people rather than replaces judgement.”.	Should read: replacing.



The capability impact, page 20.	“When regulators share what they learn, including what worked, what failed, and what surprised, then the whole regulatory system benefits.”.	Fix grammar.
3.1 Setting ethical boundaries, page 22.	“Earlier sections outlined how AI can change regulatory work, including faster processing, the need for deliberate human judgement, HITL decision-making, and the importance of testing and reversibility.”.	Rewrite for clarity.
Transparency and explainability, page 23.	• regulators should be open about when AI is involved and what role it plays • decision-makers must understand how AI outputs were generated and their limitations • affected parties should be able to understand how AI-informed decisions were reached.	Very good. Best practice may include explaining how models work.
Fairness, bias, and privacy: protecting people from unintended harm, page 23.	“Using data that is relevant and appropriate to the regulatory context”.	Hard to control this in situations where the regulator is procuring a tool from the market. May mean the agency data that is going into the model rather than training data. Unclear.
Transparency and explainability, page 23.	“Where Māori data is involved, it should be treated as taonga, with decisions grounded in the principles of Te Tiriti o Waitangi and data sovereignty expectations.”.	Very good. Should link to an external resource however. It would be good to explain how.



Transparency and explainability, page 23.	“Where AI-supported processes affect the public, flexibility is important so people are not disadvantaged by digital access, literacy or circumstance.”.	Very good.
Get clear on the ‘what’, ‘when’ and ‘why’, page 24.	“Are the benefits realistic relative to the effort, cost and maturity of a regulators practice?”.	Apostrophe missing.
Get clear on the ‘what’, ‘when’ and ‘why’, page 24.	“Has a risk assessment been completed?”.	Is there a government standard AI risk assessment protocol? If not, the Guidance should advocate for one. Otherwise, inconsistent risk assessment practices across Government is an area of risk.

6. Limitations and Disclaimer

This independent review has been prepared by LensenMcGavin AI exclusively for the Ministry for Regulation. The observations and recommendations contained in this report are based on the draft guidance document provided to us in March 2026. This report focuses on strategic, technical, and operational considerations related to AI deployment in regulatory contexts and does not constitute formal legal advice.

AI Disclosure Statement: Generative AI tools (e.g. Gemini) were used to refine writing and structure. All suggestions, critiques, and comments were produced by the authors. All GenAI output was reviewed and edited by the authors.



LensenMcGavin AI

S 9(2)(a)

S 9(2)(a)

INVOICE

10/4/2026

INVOICE NO. #107

GST NO. 114-614-313

Invoice due: on or before 20/05/2026

MINISTRY FOR REGULATION
BOWN HOUSE, LEVEL 18
WELLINGTON CENTRAL
WELLINGTON 6011

DESCRIPTION	QTY	UNIT PRICE	TOTAL
Independent Expert / Peer Review of Draft AI Guidance (per hour)	20	S 9(2)(b)(ii)	S 9(2)(b)(ii)
Subtotal(excl GST):			S 9(2)(b)(ii)
Total GST:			S 9(2)(b)(ii)
Balance Due			S 9(2)(b)(ii)

Please make payment to:
LENSENMCGAVIN AI LIMITED
S 9(2)(k)
When paying please use reference: 107

Government Model Contract-Lite

AI Guidance Peer Review – LensenMcGavin AI

The Parties

The Buyer [Ministry for Regulation](#) NZBN: 9429051927295

MINISTRY FOR REGULATION
BOWEN HOUSE, LEVEL 18
1 BOWEN STREET
WELLINGTON CENTRAL
WELLINGTON 6011

and

The Supplier LensenMcGavin AI Limited (NZBN: 9429052526190)

S 9(2)(a)

Agreement

The Supplier will provide the following goods and/or services (as applicable) to the Buyer by the date and for the price recorded in the table below (excluding GST). The price includes all expenses.

Description	Delivery date	Specific Standards (if any)	Price (exc GST)
Independent peer review of draft AI guidance for regulatory practice. This includes: reviewing the draft guidance document; providing written feedback on clarity, framing, and regulatory defensibility; and suggesting improvements to strengthen practical applicability for senior regulatory leaders.	31/03/2026	Services to be delivered with due care, skill, and diligence. Feedback to be practical, concise, and suitable for a senior regulatory audience with a non-technical or IT background or capability.	S 9(2)(b)(ii)

Invoice and payment

The Supplier will invoice the Buyer for the goods and/or services when all goods and/or services have been provided.

If delivery of the goods and/or services has been satisfactory and the Buyer receives a valid tax invoice: on or before the 5th business day of the month, the Buyer must pay that tax invoice by the 20th calendar day of that month.

Start date	This contract will start on 31/03/2026.
End date	This contract will end on the date when all goods and/or services (as applicable) have been delivered or performed to the required standard and the tax invoice has been paid by the Buyer, unless the Parties agree in writing to an earlier end date, or this contract is terminated earlier by the Buyer giving 5 business days' notice.
Standards	<u>FOR SERVICES:</u> (as applicable) The Supplier will provide the services with due care, skill and diligence and to the Specific Standards specified above (if any). <u>FOR GOODS:</u> (as applicable) All conditions and warranties under the Sale of Goods Act 1908 apply to the supply of goods under this contract. The Supplier must comply with the <i>Supplier Code of Conduct</i> issued by the Procurement Functional Leader (see procurement.govt.nz).
Resolving disputes	The Parties will use their best endeavours to resolve any dispute or difference that may arise under this contract through direct negotiation. If the Parties cannot resolve a dispute by negotiation, either Party may refer the matter to mediation. The Party requesting mediation must notify the other Party in writing. Each Party will meet their own costs of resolving the dispute.
Law	This contract will be governed and interpreted in accordance with New Zealand law.

Signed for and on behalf of the Buyer:	Signed for and on behalf of the Supplier:
S 9(2)(a)	

Name:	Alex McMinn	Name:	S 9(2)(a)
Position:	Head of Regulatory System Capability	Position:	Director
Date:	20/03/2026	Date:	23/03/2026