MINISTRY OF SOCIAL DEVELOPMENT
TE MANATŪ WHAKAHIATO ORA

Governance Guide

Model Development Lifecycle

MSD Reference: A13094818

October 2021

Contents

Introduction

Good product development governance ensures effective decision-making throughout the development, deployment and ongoing maintenance. In product development projects, it's common to have additional advisory boards alongside traditional project governance frameworks. This gives decision makers confidence that the product is maximising opportunities and responsibly managing potential risk and harms that are technical. Additional advisory boards are common in areas such as engineering or IT.

For data science products, including operational algorithms, these are less common. This is not because they're not needed, but because it's a new field and governance of data science products isn't yet well understood.

Governance specific to operational algorithms is crucial to maximise opportunities while effectively managing potential risk and harms. This is because the opportunities and risks for data science products aren't limited to the appropriate use of technical methods that could be covered with traditional frameworks like quality assurance and review. The opportunities and risks are wide reaching and span the areas of legal, ethical and te ao Māori, all of which can't reasonably be expected to be managed by a data scientist alone.

To make effective decisions in the development of an operational algorithm, decision makers need to feel assured that the technical, legal, ethical and te ao Māori opportunities and risks have been well managed.

This document provides a governance framework to give decision makers this confidence and a framework for identifying, classifying and managing risks.

Using this document

This document is a guide to apply additional supporting governance over and above any current governance, to specifically support decision-making for operational algorithm products. The target audience for this document are project managers and coordinators. It assumes the standard project management knowledge associated with those roles.

This document covers:

1. Governance framework for operational algorithms
2. Risk identification, classification and management
3. Technical Advisory Group (including Appendix 1: Technical Advisory Group Terms of Reference template).

To support the governance framework for operational algorithms there should be a documented sign-off process with approval from each of the following:

- L2 Analytics Owner
- L2 Business Owner
- L2 Communication Owner
- L2 IT Owner
- L2 Analytics Owner - Ethics
- Technical Advisory Group review outcome
- L3 Operational Algorithm Owner (authorising go-live)

Appendix 2 on page 15 contains alternative variants of the sign-off process that are based on the same framework but are suitable to different levels of algorithm risk and project sizes.

Governance framework for operational algorithms

Overview

As with any projects, those involving operational algorithms have their own project management process and governance frameworks to ensure they run smoothly, on budget, with timely deliveries, and remain within scope.

This section provides an additional governance framework to identify and reduce technical, legal, ethical and te ao Māori associated risks specific to operational algorithms and to maximise opportunities. In this context, technical refers to the accuracy and validity of the algorithms.

The purpose of this operational algorithm governance framework is to:

- ensure the right level of technical, legal, ethical and te ao Māori knowledge, expertise and experience is available at key points in the product lifecycle
- assign single points of accountability for sign-off
- ensure benefits and opportunities are maximised
- give confidence to individuals who are responsible for sign-off of deliverables that have had technical, legal and ethical risks identified and appropriately managed.

Operational algorithm governance framework

The operational algorithm governance structure includes three 'approval gates' designed to identify and reduce risk as the operational algorithm lifecycle progresses. Deliverables must be approved for sign-off at each gate and risks assigned to that level (see Risk identification, classification and management) must be owned to progress to the next phase.

Figure 1: Reducing risk through clear accountable sign-off.

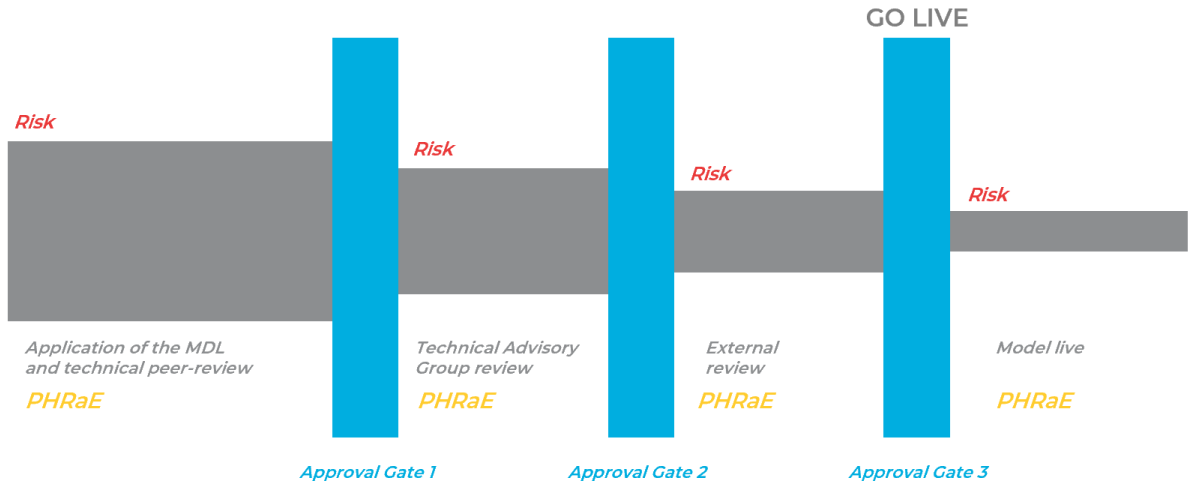

The role of the PHRaE

An organisation's Privacy, Human Rights and Ethics (PHRaE) framework is a key tool for all operational algorithms required throughout all stages of this operational algorithm governance framework. It is reflected in detail in the *User Guide*, and in the sign-off process for operational algorithms.

For Approval Gate 1, a draft version of the PHRaE should be completed. For Approval Gate 2 and Approval Gate 3, an individual or individuals must be assigned as a Level 2 and Level 3 approver, respectively. For these roles, this governance framework also provides sign-off sheets.

In practice, if the external review suggests significant changes the level 2 sign-off may be repeated before and after the external review.

Sign-off sheets

Sign-off sheets are available as part of this governance framework, and make clear the following:

1. What is required to be signed off.
2. What decisions the approver can make in relation to sign off.
3. The benefits and opportunities.
4. What risks must be owned.
5. What risks must be escalated for ownership.
6. What information must be noted.

Table 1: Sign-off documentation for operational algorithms.

Relevant section of MDL Data Science Guide for Operations	Approval gate level	Sign-off
Data science methods in operation	2	L2 Analytics Owner
Data science integration: working with the business unit receiving your service	2	L2 Business Owner
Data science integration: working with business units who provide support	2	L2 Communication Owner
Data science integration: working with business units who provide support	2	L2 IT Owner
Data science application of principles and frameworks to manage potential risks and harms	2	L2 Analytics owner - Ethics
All	NA	Technical Advisory Group review outcome
All	3	L3 Operational Algorithm Owner

Setting up operational algorithm governance

The operational algorithm governance framework should be set up at the same time as the organisation's usual project management process and project governance framework. It is recommended that this happen before product development officially starts. This includes assigning clear sign-off roles (Level 2 and 3) and the Technical Advisory Group.

Level 1 approval should be a team's 'usual' method of quality assurance. For example, in a data science team there is usually a principal or technical lead who provides quality assurance. No sign-off sheets have been provided for level 1 approval.

Maximising benefits and opportunities

Introduction

Maximising benefits of operational algorithms, and clearly articulating these is crucial to understand how these balance the potential risks and harms. Decision makers need to weigh up the benefits against the risks to effectively determine whether the product should be deployed into operation.

Roles and responsibilities

Everyone working on operational algorithms has a responsibility to ensure benefits and opportunities are maximised. There are two key roles that facilitate accountability.

The role of Technical Advisory Group

The Technical Advisory Group has a key role to play in ensuring that the stated benefits are realistic and that those benefits are maximised. They can do this by providing guidance, asking the right questions and identifying opportunities to maximise the benefits.

The Technical Advisory Group may also be responsible for identifying when stakeholder engagement is required, and ensuring it is used in a way to maximise benefits and opportunities. This is separate from co-design, which is triggered through the PHRaE.

In practice, stakeholder engagement may be triggered in any stage or by any manager (business owner, analytics owner, communications owner, DCE, and so on). The role of the Technical Advisory Group is to ensure this has happened and if not, to recommend it does happen.

The PHRaE and the Guidelines for Trusted Data Use[1] are both useful tools to guide discussions in stakeholder engagement.

These requirements are captured in the Technical Advisory Group's Terms of Reference (ToR) on page 12.

The role of the analytics owner

The key tasks associated with this role are to:

- empower their staff to engage with the rest of their business to discover new ideas
- be active in identifying benefits and opportunities
- provide guidance and ask the right questions.

Examples of benefits and opportunities

The benefits may be to the wider organisation, frontline staff, service providers or service users. Examples include the following:

- Faster decision-making. This has benefits for both service providers, who may be able to proceed immediately without a follow up appointment, and service users, who can access services faster and proceed with greater surety of service access.
- Cost savings via automation of simple tasks. This has financial benefits and frontline staff can focus on tasks that make better use of their skills and create more value for service users.

[1] https://www.aisp.upenn.edu/wp-content/uploads/2019/08/Trusted-Data-Use_2017.pdf

- More consistent decision-making. This has benefits for everyone involved. For service users decision-making can become more standard and fair, for staff designing or delivering the service the requirement to manage variability in decision-making can reduce, and for any users of operational data it can become cleaner or less variable (this has many benefits to future data projects and products).
- More transparent decision-making. This also has benefits for everyone involved. This metric will help decide the scale of the opportunity and the extent of desired change.

Risk identification, classification and management

Overview

This section is repeated in the *Data Science Guide for Operations* as it is necessary for all audiences and team members.

Risk management **does not** eliminate risk. Risks are never closed[2] – they exist as long as we undertake an activity to achieve an outcome. We can, however, influence the level of risk in terms of its potential impact or likelihood. The purpose of risk management is to enable informed decisions. It allows us to consciously choose whether we accept a given level of risk or act to reduce that risk.

To ensure that risk management is consistently understood and applied across MDL products, a common approach is needed. This section provides a simple approach to identifying, classifying and managing risk for MDL products.

This doesn't replace risk management policies in your organisation. Instead, it supplements current best practice with a primary focus on the risks associated with operational algorithm projects.

Roles and responsibilities

Everyone working on operational algorithms has a responsibility to ensure risks are identified and managed. While the ownership of risk resides with accountable individuals, all team members involved have a role to play.

For data scientists, the *Data Science Guide for Operations* is a practical guide to avoid risks associated with technical error. These risks need to be identified so they can be managed and owned.

Beyond technical error, there are other risks that can't be avoided that need to be identified and categorised so they can be managed and owned.

The role of the data scientist

The key tasks associated with this role are to:

- use the MDL as a first barrier to avoid risk
- ensure identified risks are classified, recorded and escalated appropriately based on the governance framework for the product.

The role of the analytics owner

The key tasks associated with this role are to:

[2] Issues are risks that have already occurred. Issues will usually be dealt with as a project progresses and once resolved can be closed. Any issues that remain outstanding can be thought of as a type of risk. In those cases, the mitigations take the form of planned remedial actions.

- ensure appropriate governance is set up for operational algorithm products
- actively work to identify and manage risks, including identifying owners for the controls mitigating or reducing any impacts
- ensure identified risks are classified, recorded and escalated appropriately based on the governance framework for the product
- clearly describe the residual risk for any risks that have not been fully mitigated.

The role of Technical Advisory Group

As in the maximising benefits and opportunities section above, the Technical Advisory Group has a key role to play in ensuring risks and potential harms are identified and appropriately managed.

The group is also responsible for identifying when stakeholder engagement is required, and ensuring it is used in a way that appropriately identifies and manages risks and harms.

Main ideas

Risks must be categorised

When risks are identified, they must be recorded in a risk register (see ToR on page 12) and classified on identification.

Figure 2 provides a risk classification matrix. This is a self-assessment tool.

Impact is the severity of the impact should the risk occur. **Likelihood of occurring** is how likely the risk is to occur.

Figure 2: Risk classification matrix.

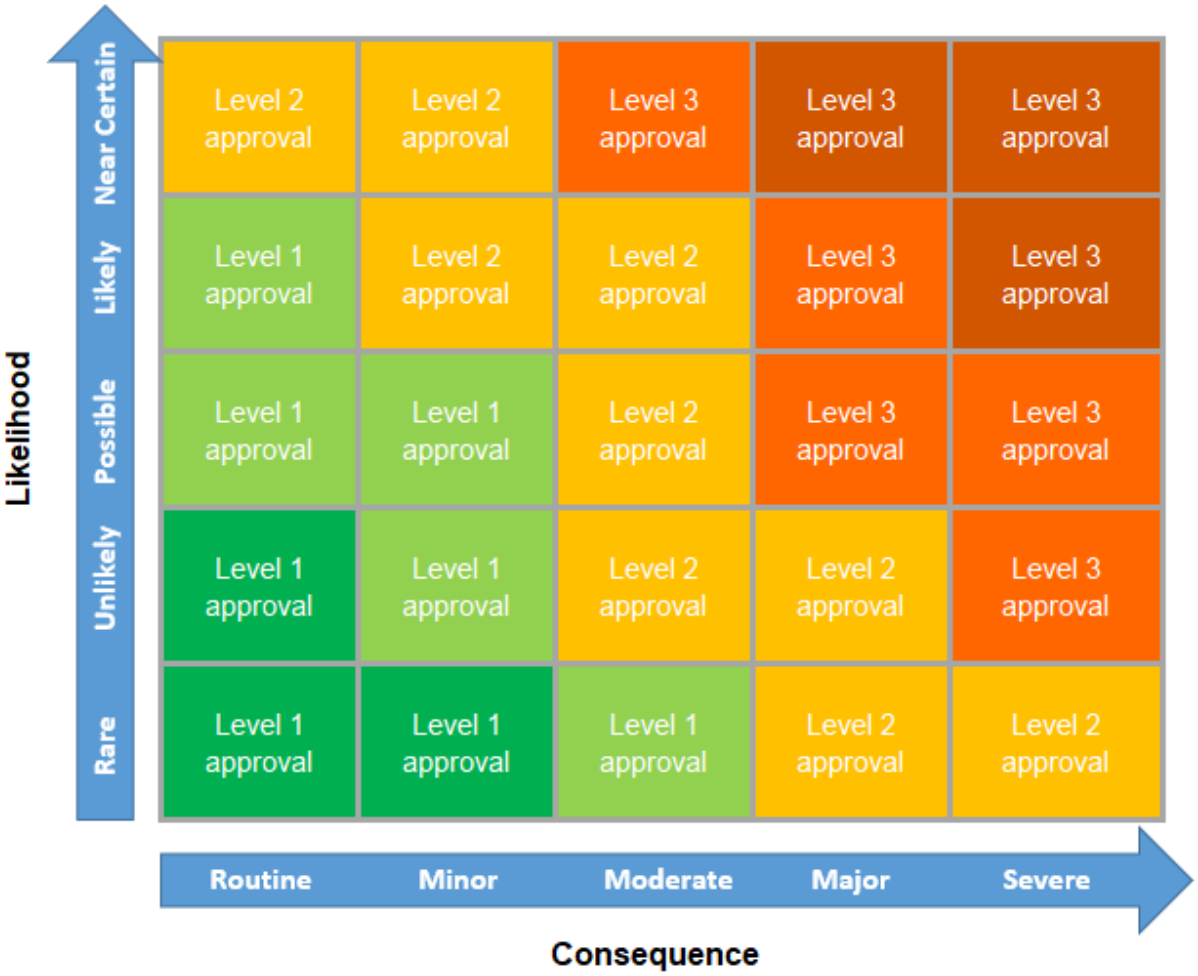

Risk impact

Examples of low risks could include:

- Working with highly skewed or missing data that results in unequal sample sizes and variety. There are technical risks associated with using such data, so it needs to be pre-processed before the application of analytical models – at the same time ensuring the pre-processing doesn't affect the business context of the data.
- Low sample sizes or small effect sizes, or other issues that affect the power of the statistical analysis. These are quite common risks associated with analytical models and require a power analysis to determine what can be realistically achieved from the available data.
- Model tuning could be resource-intensive and take a huge amount of processing time to arrive at a suitable level of calibration. This is dependent on the complexity of the model itself, and the number of parameters to be tuned. Each additional parameter increases time or resources needed for tuning, and sub-par tuning increases the risk of unpredictable model performance.
- The business loses the staff skills required to make the decisions the algorithm now makes. This will expose the organisation to risks should the algorithm need to be turned off (eg due to a legal decision). These risks may not have been considered when the algorithm was developed and subsequently freed up frontline staff who had experience making those decisions to work in other areas.

Examples of medium risks could include:

- Having to make compromises on model accuracy to ensure interpretability (or vice versa). There are chances that an easily interpretable model can have lower accuracy than a more sophisticated, black-box analytical model. These are often contradictory goals that could impact model outcomes.
- Model overfitting – repeated testing and tuning of models – can easily propagate unconscious biases regarding data, especially if the same datasets get re-used. The impact might be unexpectedly low model performance when the model is used for scoring real-world data and may require extensive retraining.
- Noisy data or influential outliers that violate key technical assumptions in a statistical model leading to unexpected model behaviour or poor performance in model metrics.
- The business uses the predictions of the algorithm in a new way that was not considered when the PHRaE was completed. This could result in unethical uses and uses of the data for which the organisation does not have appropriate permission for.

Examples of high risks could include:

- Developing analytical models that were trained on examples not entirely reflective of real-world scenarios in terms of business context or data quality. This could lead to unexpected poor model performance and require significant re-engineering efforts.
- The use of data for purposes where consent or social licence isn’t clear. These could be concerns relating to privacy or ownership of the data, or the use of the data in a context that was not agreed to by its providers. The impact might be that the results from this project may have to be completely redacted, even if the study was well-designed and the results were insightful. There might also be significant legal and reputational impacts for the organisation.
- Use of a variable with a large predictive power, but with inappropriate or questionable meaning in an analytical model. For instance, the use of ethnicity or gender identity might be particularly concerning in certain business contexts but may have predictive power because it acts as a proxy for certain unobservable characteristics. This is quite a common occurrence but will have a significant social impact by propagating biases and a reputational impact for the organisation.

- That frontline staff become overly reliant on the predictions of the model. This is relevant to augmented decision-making and would mean that the model and human decision-making are no longer working as designed. This could have a significant impact on the quality of decisions as well as on the ethics of those decisions, which may rely on human decision-making as a safeguard in certain situations.

Risk likelihood

Table 2 provides guidance on selecting the appropriate likelihood of occurrence level to categorise risk.

Table 2: Risk likelihood of occurring guide.

Likelihood of occurring level	Under what circumstances could the risk occur	When is the risk expected to occur	What controls are in place to prevent the risk occurring	Has the risk occurred before
High	Probable: Likely to occur often during standard operations.	The risk is expected to occur within the next 6 – 12 months.	No effective controls or weak controls, eg limited business controls, with no audits performed.	Has happened in the past and no compensating controls have been implemented.
Medium	Occasional: Likely to occur sometime during standard operations.	The risk is expected to occur within the next 1 to 3 years, with a 20%-50% expectation that the risk will occur during the next 12 months.	Minimal controls, eg some business controls, with some audits performed.	The risk has occurred in other similar organisations with similar levels of controls in place.
Low	Improbable: Unlikely to occur or is only expected to occur in exceptional circumstances, such as deliberate fraud or activity beyond control of business actions.	The risk is not expected to occur within the next 5 years, and there is a less than 20% expectation that the risk will occur during the next 12 months	Effective controls, eg timely business controls, with internal & external audits performed.	The risk hasn't occurred in the business or in other similar organisations.

Risks must be owned and treated

Once risks have been identified and categorised, the controls to mitigate them need to be owned. Ownership depends on the category of risk and is assigned based on current risk.

Raw risks are risks that have not yet been treated. Current risks are risks that have been treated. The risk register should capture the category of the risk before treatment/mitigation (raw), and after (current).

In the event the risk moves category after it has been assigned to a control owner, it may be reassigned to a new control owner to reflect the new level of risk.

Figure 3: Risk classification matrix impact of treatment.

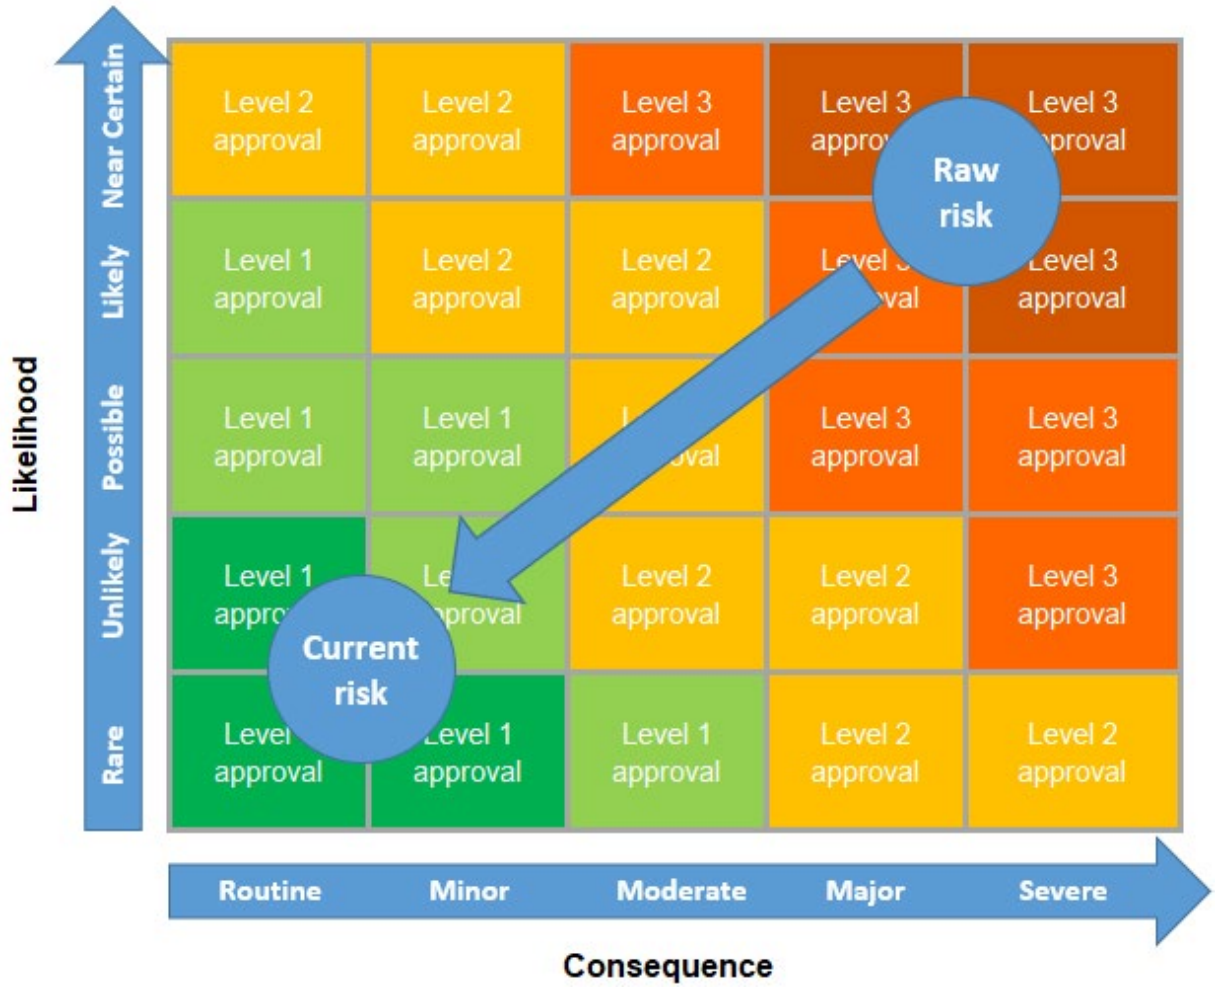

Mitigations and controls need owners too

The responsibility of the person owning a risk is to be confident that the treatments applied to mitigate that risk are effective and will continue to be so. Most of the time that person will not be the one who manages and implements the control itself.

Identifying control owners as part of the risk register provides clarity for the person signing off on a risk and is a way for the control owner to confirm the control is an appropriate fit for the risk it is applied to.

Owners of controls should be able to demonstrate that the control is effective, and that thought has been given to the maintenance and monitoring of the control.

One control sometimes mitigates more than one risk. For example, IT security around firewalls speaks to both the robustness of the delivery system and is an important means of protecting privacy. It can be helpful to build a library of well understood controls and the types of risk they commonly apply to.

Technical Advisory Group

Overview

The Technical Advisory Group has three responsibilities for operational algorithm products:

1. ***Guidance*** - providing guidance to the project team throughout the operational algorithm product lifecycle.
2. ***Review*** - review before Level 2 sign-off (the outcome of the review is to approve or endorse).
3. Recommendation:
    a. recommending (or not) external review
    b. recommending (or not) stakeholder engagement

The Terms of Reference for the Technical Advisory Group are in Appendix 1 on page 12. The ToR provides details of how these three responsibilities are ensured through:

- decision-making
- benefits and opportunities
- risks
- stakeholder engagement
- ethical review
- legal review
- te ao Māori review
- membership
- skills
- administration
- recommending external review.

Guidance

The Technical Advisory Group is responsible for providing guidance throughout an operational algorithm project. This means any member of the project team can contact members of the advisory group at any time to ask questions.

Review

The Technical Advisory Group can approve OR endorse deliverables in an operational algorithm project, so these can progress to the next phase (see Figure 1 on page 2). The decision-making power needs to be decided when the Technical Advisory Group is set up.

It is recommended the Technical Advisory Group only endorse the deliverables to give the Level 2 approver confidence in the algorithm's development. If the Level 2 approver is also a member of the Technical Advisory Group, the group may instead be able to approve deliverables. That is, the group will function to review AND approve deliverables.

Recommendation

External review

The Technical Advisory Group is responsible for recommending whether any aspect of the operational algorithm needs to be sent externally for review.

There are three types of review the Technical Advisory Group can choose from:

- Technical review
- Legal review

- Ethical review

Where external review is required, the Technical Advisory Group must specify:

- which type/s of review is required
- what aspect of the product needs to be reviewed (the purpose of the review)
- who the review will be completed by.

Stakeholder engagement

In practice, stakeholder engagement (or consultation) may be triggered at any stage or by any manager (business owner, analytics owner, communications owner, DCE, and so on).

The role of the Technical Advisory Group is to identify when stakeholder engagement is required, and ensure it is used in a way that maximises benefits and opportunities and identifies and manages potential risks and harms.

The PHRaE and the Guidelines for Trusted Data Use[3] are both useful tools to guide discussions in stakeholder engagement.

Stakeholders can include, but are not limited to:

- clients
- advocates
- service users
- affected parties
- other organisations.

Setting up the Technical Advisory Group

The Technical Advisory Group should be set up at the same time as the operational algorithm governance framework.

The Technical Advisory Group is not a stakeholder group

Stakeholder groups should have representation from all affected peoples, including operational business units and the organisation's customers. Feedback from stakeholder groups should be available to the Technical Advisory group to support its discussions.

Assigning a Technical Advisory Group chair

The Technical Advisory Group must have a chair. The chair is responsible for:

- chairing Technical Advisory Group meetings
- casting a deciding vote when a consensus can't be reached between Technical Advisory Group members.

To avoid any conflict of interest the chair should not be directly involved in completing work on an operational algorithm product or be directly responsible for the team who is (for example, the Analytics Owner or L3 sign-off manager).

In this situation, if this is the only appropriate chair, the chair will not have a deciding vote in the event a consensus can't be reached. Instead, an unequal number of people should make up the Technical Advisory Group to avoid any possible draws.

[3] https://www.aisp.upenn.edu/wp-content/uploads/2019/08/Trusted-Data-Use_2017.pdf

Appendix 1: Technical Advisory Group Terms of Reference template

[This text is guidance text to be used when setting up advisory group]

The Technical Advisory Group has three responsibilities for operational algorithms:

1. ***Guidance*** - providing guidance to the project team throughout the operational algorithm product lifecycle
2. ***Review*** - review before Level 2 sign-off (the outcome of the review is to approve or endorse)
3. Recommendation:
    a. recommending (or not) external review
    b. recommending (or not) stakeholder engagement

Decision-making

4. Decisions are recorded in a decision register, which is stored at the following location *[add document location here]* for access by all members of the operational algorithm project team.
5. Where possible, decisions are made by consensus. If consensus can't be reached, the final decision rests with the majority, with the chair exercising a deciding vote if required.
6. There are four possible decision outcomes from the advisory group:
7. Endorsed
8. Endorsed with conditions
9. Endorsed with minor changes
10. Not endorsed

Benefits and opportunities

11. The Technical Advisory Group actively works to maximise benefits and opportunities.

Risks

12. Risks are recorded in a risk register, which is stored at the following location *[add document location here]* for access by all members of the product development team.
13. The Technical Advisory Group actively works to identify and appropriately manage risks.
14. The Technical Advisory Group classifies risks and ensures risks are escalated to the correct level.

Stakeholder engagement

15. The Technical Advisory Group will identify when stakeholder engagement is required, and ensure it is used in a way that maximises benefits and opportunities and identifies and appropriately manages risks and harms.

Ethical review

16. Ethical review involves considering and balancing:

17. the potential for harm (physical, emotional or cultural) to any participant and ways to minimise this harm
18. the potential value and benefits of the operational algorithm for customers.
19. The Technical Advisory Group is guided by ethical standards that are based on those of the National Ethics Advisory Committee for health and disability research in New Zealand, and guided by the *[insert organisation name]* Privacy, Human Rights and Ethics framework for the safe use of data, and Statistics New Zealand's Principles for the Safe and Effective Use of Data and Analytics.

Legal review

20. The Technical Advisory Group is guided by the three principles of te Tiriti o Waitangi/the Treaty of Waitangi – Participation, Protection, and Partnership.
21. The Technical Advisory Group is guided by legislation relevant to *[insert organisation name]*, Human Rights Act and the Privacy Act.

Te Ao Māori review

22. The Technical Advisory Group is guided by te ao Māori principles, such as principles of Māori data sovereignty.

Membership

23. The Technical Advisory Group will be set up to ensure that it, as a whole, has the skills, knowledge and ability to fulfil its purpose. The Technical Advisory Group size is dependent on the size of the associated opportunities and risks.
24. The Technical Advisory Group can be made up of internal and/or external members.

Skills

25. The essential skills required by the Technical Advisory Group as a whole include:
26. ethical expertise
27. legal expertise
28. expertise in the application of te ao Māori principles
29. data science technical expertise
30. business operational knowledge (subject matter experts).
31. A member from the PHRaE team.

Administration

32. Fortnightly meetings should be set up prior to the product development starting whilst urgent matters can be dealt with via email [Include details of meeting frequency and availability outside of meeting times to project members. To be completed as part of MDL governance project set up].
33. Other people involved in the project may attend the Technical Advisory Group by invitation of the group, including a group administrator.
34. The Technical Advisory Group has no financial delegation.

Recommending external review

35. There are three types of review the Technical Advisory Group can choose from:
    - Technical review
    - Legal review
    - Ethical review

36. In the case where external review is required, the Technical Advisory Group must specify:
    - which type/s of review is required
    - what aspect of the product needs to be reviewed (the purpose of the review)
    - who the review will be completed by.

Appendix 2: Sign-off and governance variants

Different degrees of sign-off and governance will be appropriate to different projects based on their size and associated risk. Different governance processes are outlined below, along with when to use them and their advantages and disadvantages.

The governance process outlined in the rest of the *Governance Guide* is called 'Full governance with distributed sign-off' and is appropriate for large projects with medium to high levels of risk.

Organisations should choose one type of governance that is suitable to the highest level of risk that will be associated with their algorithms. Large government organisations should be using a 'Full Governance' approach but may use 'Reduced' or 'Minimum Governance' as a steppingstone to 'Full Governance'. This steppingstone should be short-lived and only address low-risk algorithms. Its goal is to allow the organisation to learn about governance and identify where they should add people to their Technical Advisory Group.

Small government organisations can use 'Reduced' or 'Minimum Governance' as their permanent solution, but they will be limited to doing small projects with low risk.

The Technical Advisory Group should always include someone with a good knowledge of te Tiriti o Waitangi/the Treaty of Waitangi and representatives of groups that are often disadvantaged by algorithms.

Variant 1: Full governance with accumulated sign-off

When to use it: Large projects with medium to high risk.

Details:

- The same as 'Full governance with distributed sign-off' except that the L2 owners don't own any outstanding risks.
- Instead the L2 owners can only endorse the acceptance of the parts of the algorithm they are responsible for.
- All the sign-off, as well as ownership of all outstanding risks, lies with the L3 owner.

Advantages:

- All risks are still identified and mitigated in a distributed way.
- All risks are owned by a single person, which is a simpler approach.

Disadvantages:

- One person must evaluate all the risks and opportunities. This requires them to understand a lot of different parts of the algorithm and the associated project in detail.
- All the sign-off occurs in one go at the end of the project by someone who isn't involved in the details of the project. There is a chance that they will find the accumulated risk too great and not sign off the algorithm.

Variant 2: Reduced Governance

When to use it: Small to medium sized projects with low to medium risk.

Details:

- Technical Advisory Group may be smaller (2 to 3 members) or meet less often. It will focus on identifying and mitigating risk rather than improving the performance of the algorithm.
- Less chance of external review or less stakeholder engagement required.
- Sign-off documentation remains the same as for 'Full Governance'.
- Several of the L2 owners may be the same person.
- Technical Advisory Group should be based on at least one expert on governance and algorithms.

Advantages:

- Faster to setup and less overhead.
- Allows the organisation to learn what they require from their 'Full Governance' approach and Technical Advisory Board.
- Identifies risks and possible mitigations at low cost.

Disadvantages:

- Only applicable to lower risk projects.
- Technical Advisory Group will have less opportunity to ask questions or make suggestions that may improve the algorithm.
- There is a chance that risks may be missed because the Technical Advisory Group only contains a small number of viewpoints.

Variant 3: Minimum Governance

When to use it: Small projects with low risk.

Details:

- Technical Advisory Group may be smaller (2 members) or meet less often. It will focus on identifying and mitigating risk rather than improving the performance of the algorithm.
- Less chance of external review or less stakeholder engagement required.
- Single sign-off document.
- The L3 and L2 owners may be the same person.

Advantages:

- Faster to setup and less overhead.
- Allows the organisation to learn what they require from their 'Full Governance' approach and Technical Advisory Group.
- Identifies risks and possible mitigations at low cost.

Disadvantages:

- Only trivial, low-impact algorithms should be considered.

- Technical Advisory Group will have less opportunity to ask questions or make suggestions that may improve the algorithm.
- There is a chance that risks may be missed because the Technical Advisory Group only contains a small number of viewpoints.