AWS Well-Architected Framework

Māori Data Lens

Māori Data Lens : AWS Well-Architected Framework

Table of Contents

Māori Data Lens - AWS Well-Architected Framework

Publication date: **August 1, 2024** (*Document history*)

This paper describes the Māori Data Lens for the AWS Well-Architected Framework. The purpose of the Māori Data Lens is to support customers who want to apply Māori data considerations when using AWS services and the Well-Architected Framework. This lens includes general design principles, best practices, and specific guidance. It offers a way to understand AWS best practices and how these may be considered in relation to Māori data. By using these best practices, you can consider how to design and operate reliable, secure, efficient, cost-effective, and sustainable workloads in the cloud.

Abstract

At AWS, we understand that data is our customer's most treasured possession. We offer the most comprehensive set of services, tools, and expertise to help customers protect their data, underpinned by a flexible and secure cloud computing environment.

The AWS Well-Architected Framework helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications. The Framework provides a consistent approach for customers and AWS Partners to evaluate architectures, and provides guidance to implement designs that scale with your application needs over time.

Foreword

Tēnā koutou katoa,

Warm greetings to all those considering how to store and use data in Aotearoa. We developed this lens to give you a place to start.

We gathered a group of Māori technology and data experts to advise on practical considerations you or your company may want to think about. The group consists of the following participants:

- Atawhai Tibble
- Renata Hakiwai
- Ngapera Riley
- Lee Timutimu

- Wade Reweti
- Nikora Ngaropo
- Eli Pohio

In this rapidly evolving digital landscape, the importance of understanding, organising, and harnessing data cannot be overstated. For our Māori communities, this holds even greater significance, as data can be a taonga or a treasure that represents the collective wisdom and knowledge passed down through generations.

The Māori Data Lens is a testament to our commitment to preserving and elevating our cultural heritage, as well as fostering innovation in the digital age. This lens, designed with input from our communities and shaped by the insights of Māori tech experts who are practitioners, provides a structured approach to data management that aligns with our values, traditions, and aspirations and creates a synergy with the AWS Well-Architected Framework pillars.

These ideas are not mandates but rather thoughtful considerations between the intricate balance of technology and tradition. In our collective journey toward data-driven decisions and actions, we must ensure that the lens remains a living embodiment of our cultural identity and wisdom. It is not just a tool, but a reflection of our unique worldview, a guardian of our knowledge, and a bridge between past, present, and future.

We extend our gratitude to all those who have contributed to the development of this lens, and we encourage everyone, Māori and non-Māori alike, to engage with it in a spirit of ongoing collaboration and respect. Together, we can weave a tapestry of data that empowers our communities, enriches our understanding, and ensures that the legacy of Aotearoa endures for generations to come.

Nā tō rourou, nā taku rourou, ka ora ai te iwi.

With your food basket and my food basket, the people will thrive.

Mauri ora!

Overview

AWS has helped New Zealand organisations innovate, succeed, and grow globally since 2013. There are thousands of active customers using AWS every month. Some customers have asked us how to use AWS in support of indigenous data, and specifically Māori data. They have also asked what considerations they can make to support Māori digital aspirations and interests. These interests include protecting data that is considered taonga, or a treasured possession. In other cultures, this can mean sacred data. Other customer interests include how data is made available to Māori and how tools are developed, so that they can use their data to deliver better outcomes to themselves, their communities, and Aotearoa New Zealand.

We developed this document together with Māori advisers as a practical resource for customers working with Māori data, recognising the importance of Te Tiriti o Waitangi, the Treaty of Waitangi. This is the first published AWS lens that focuses on indigenous data. This document is intended for those in technology and data roles, such as chief technology officers (CTOs), chief information security officers (CSOs/CISOs), architects, developers, and operations managers.

Pillars of the Well-Architected Framework

The Well-Architected Framework consists of six pillars: operational excellence, security, reliability, performance efficiency, cost optimisation and sustainability. The Framework describes key concepts, design principles, and architectural best practices for designing and running workloads in the cloud. The following sections provide an overview of each pillar, important design principles, and other considerations for your organisation as it relates to Māori data.

Purpose

This lens is not a comprehensive guide, but instead is based on AWS and customer best practices, combined with Māori data considerations. It is not intended to be an authoritative, prescriptive checklist. Instead, it provides a technology-focused perspective that centres around designing, building, and operating technology solutions. This document does not replace the value of working with cultural advisors that have cultural knowledge and specific context of the data. This document was developed with Māori advisers that have a deep understanding of Māori data interests, but it should not replace other consultations with your Māori customers that have relevant cultural context of the data or their specific tikanga (protocol) over that data.

An AWS Well-Architected Framework lens is designed to be a specific way of looking at or applying the AWS Well Architected Framework and its six pillars. By using the lens, you can learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable applications and software in the cloud. It provides a way to consistently measure your cloud environment against best practices and identify areas for improvement. Applying Well-Architected Framework best practices to your workload increases the likelihood of business success and protection of cultural and historical knowledge. We developed specific domain lenses, such as the Government Lens, the Healthcare Lens, or the Financial Services Industry Lens.

Guidance and lenses are intended to help AWS customers and other organisations reflect on their data policies and principles in the context of the industry or sector they operate in. This document outlines general design principles, best practices, and specific guidance for the pillars of the Framework. It offers a way to understand AWS best practices and how these may be considered in relation to important Māori data considerations. You can apply this lens to a workload in isolation. However, we recommend applying this lens in addition to a Well-Architected Framework Review to fully evaluate your workload.

AWS bringing infrastructure to Aotearoa New Zealand

On 22 September 2021, AWS announced that the new AWS Asia Pacific (Auckland) Region. The Region adds to the continuing AWS investment in and commitment to Aotearoa New Zealand, and the long-term potential for New Zealand to be a leader in the global digital economy. This region will have three Availability Zones and can provide local customers with a preference to securely store their data in New Zealand, as well as provide even lower latency to users. We hope customers can take advantage of this significant investment in New Zealand and the relevant data controls to meet the needs of our Māori customers with specific data preferences. It can also support organisations to leverage advanced technologies such as artificial intelligence including generative AI, machine learning, internet of things (IoT), and mobile services to drive innovation. Customers can start building out their workloads in an existing Region now, and work with AWS or a trusted partner to get architectural and implementation guidance that facilitates a transition to the Auckland Region in the future.

Additional context

The first step in your Well-Architected Framework Review is deciding if you want to separately identify Māori data from other data. From there, you can decide on a classification system to help you identify what Māori data you already have or are intending to collect. To do this, your

organisation needs to decide on a robust definition for what is Māori data, as well as a mechanism to help you potentially classify, within the context of your organisation, what is and what isn't Māori data. Your organisation can also decide if it is appropriate or not for you to start identifying data as Māori and non-Māori, and what are the specific legal considerations you might wish to take. You can speak with Māori data specialists about this. After you have decided to identify data separately as Māori data, you may choose to decide on a data classification system to drive your technology choices, data governance frameworks, and security best practices over different types of Māori data. For more information, see the Security Pillar.

The definitions provided in this guidance are just a starting point. Other organisations may publish Māori data classification guidelines that you could also use as a reference. To get you started, we suggest you look at this from some practical angles:

- Has the data been provided by someone who has identified themselves as Māori?
- Will this data be distinctly useful for Māori because it has to do with their community, environment or wellbeing?
- Does this data relate to a Māori individual or community that holds value to that individual or community?

This list of considerations is not exhaustive, and these questions may take time for your organisation to answer. We strongly encourage you to speak with Māori data experts about your specific context and about what data you have collected or intend to collect. This discovery process can help you set up appropriate mechanisms for stewarding Māori data. At the heart of this is the importance of the Treaty of Waitangi (Te Tiriti o Waitangi), signed between the British Crown and more than 500 Māori chiefs in 1840. The New Zealand Government's Cloud First policy acknowledges the Treaty and has Te Tiriti-based principles for government agencies using public cloud.

This lens is just one tool for supporting customers in considering good practices. This lens should not be regarded as static or a stand-alone solution, and does not seek to convey that there is any single best approach for working with Māori data. AWS reviews and updates guidance and content continually as best practices evolve over time and we welcome feedback on this document for future revisions. Where appropriate, this guidance also includes wider organisational aspects technology and data teams can also consider.

Definitions

Builders, architects, chief information/privacy officers (CIOs/CPOs), and technologists can use the following definitions as a starting point to understand commonly used Māori terms and phrases relating to protocol (tikanga), principles, and values from a Māori viewpoint (te ao Māori).

- **Hapū:** A collection of whānau with a shared genealogy at a smaller unit than an Iwi. A collective of hapū form an iwi.
- **Iwi:** A Māori community or people made up of more than one hapū.
- **Kaitiakitanga:** Guardianship or management of both the seen and unseen worlds – this could be extended to the digital realm where past, present, and future knowledge, culture and history is increasingly being created, shared, and preserved.
- **Kotahitanga:** Unity or solidarity. The state or circumstances of being one.
- **Manaakitanga**: Hospitality, kindness, generosity, support - the process of showing respect, generosity and care for others. This can be demonstrated in how a service is delivered to someone, their family (whānau) and communities.
- **Māori data (digital):** A digital record from or about Māori, the places Māori have a connection with based on their heritage. Digitised data about iwi, hapū or Māori organisation, language, culture, resources, environments or knowledge systems.
- **Māori owned-businesses:** There is not one single agreed definition on *Māori owned* or a *Māori business*. The following combines a few suggestions for how your organisation can determine and define Māori owned-businesses for the purposes of partnership and understanding your Māori customers:
  - **Ownership:** Some define this as having at least 50% of each class of partnership interest owned by Māori people but this can be limited to certain situations. Alternatively, you can determine Māori ownership if it is owned or part-owned by a person or people where their whakapapa is Māori (have a verifiable Māori genealogical lineage) and a representative of the business identifies it as a Māori business.
  - **Self-identification:** If people promote themselves as a Māori business.
  - **Employment:** If business employs a large percentage of Māori staff.
  - **Values:** They run their business according to potentially both traditional and contemporary aspects of Māori culture and values. Examples of this include employing whānau, welcoming visitors, and using traditional practices.
- **Mātauranga Māori:** Māori knowledge.

- **Tangata whenua:** (Māori) people of the land.
- **Whakapapa:** A line of decent from one's ancestry. This places oneself in a wider context, and links oneself to land and tribal groupings.
- **Whānau:** An extended family or community of related families who may live together in the same area.
- **Whanaungatanga:** Close connection between people. It could mean kinship or a sense of family connection.
- **Workload:** A workload is a collection of resources and code that delivers business value, such as a customer-facing application or a backend process. A workload might consist of a subset of resources in a single AWS account or be a collection of multiple resources spanning multiple AWS accounts.

Te Ao Māori principles

This section explains important Māori values and principles that Māori technologists and data experts identified. Consider these design principles and the AWS Well-Architected Framework pillars when creating a culturally-responsive Māori data framework.

These considerations explain what matters to Māori and why. Use these considerations as you start designing your systems. Work with your Māori customers to understand the tikanga (the specific customary practices and protocols) to consider throughout the development process.

For example, customers may wish to consider the following high level Māori principles when working with Māori data:

- Kaitiakitanga
- Kotahitanga
- Manaakitanga
- Rangatiratanga
- Whanaungatanga
- Whakapapa

This list is not exhaustive. Consider other principles that relate to your organisation.

Kaitiakitanga

People are increasingly digitising, sharing, and preserving knowledge, culture, and history. This principle is about the care and protection of this digital taonga (treasure) for generations to come. People entrusted with Māori data can learn different options for preserving this important digital taonga. Kaitiakitanga supports the other principles and puts emphasis on the importance of protecting and preserving that data for future generations. Use this principle to guide decisions on how to collect, store, use, and protect data.

Kotahitanga

Kotahitanga describes the unity of a group of people and the ability to move together toward a shared goal. From a data perspective, you can apply this principle to understand insights and make informed decisions that benefit the collective.

Manaakitanga

This principle is closely related to kaitiakitanga. Manaakitanga is part of the actions that make up kaitiakitanga. This includes the process for engagement, hospitality, and creating the experience so that Māori genuinely know they are welcome. Using this principle includes taking a collaborative approach, developing solutions alongside partners, and potentially thinking about what this means from a Māori perspective. Applying manaakitanga may include developing a technology solution that can be used across multiple devices, enhancing accessibility through different formats and functionalities.

Rangatiratanga

Rangatiratanga is the ability to have sovereignty, leadership, and autonomy to make decisions and determine the direction you take. For example, an individual or the community has the necessary data-driven tools to determine the right journey and make decisions.

Whanaungatanga

Whanaungatanga is about relationships. In the context of the AWS Well-Architected Framework pillars, this means the relationships Māori have with data. This includes but is not limited to connecting, maintaining, sharing, and growing data. You may want to consider knowledge development, interoperability of application programming, interfaces (APIs), and open data usage.

Whakapapa

Whakapapa relates to genealogy and where we come from. For many Māori, connection to their iwi, hapū, and whānau is important. The connection is ongoing and connects people to their ancestry, ancestral history, culture, and the understanding of who they are as a people in the past, present, and future. Tracking data, where it originates from, and where it was accessed could reflect this principle.

These principles are important te ao Māori principles and we encourage users of this lens to learn more. For more detail, see Te Ara Encyclopedia of New Zealand or in printed or digital copies of Tikanga Māori: Living by Māori values by Hirini Moko Mead.

Operational excellence

The Operational Excellence Pillar provides guidance on running and monitoring systems to deliver customer value and continually improving the supporting processes and procedures. The pillar also provides an overview of design principles, best practices, and questions. You can find prescriptive guidance on implementation in the Operational Excellence Pillar whitepaper.

At AWS, operational excellence is defined as a commitment to build software correctly while consistently delivering a great customer experience. It contains best practices for organising your team, designing your workload, operating it at scale, and evolving it over time. Operational excellence focuses your team more on building new features that benefit customers and less on maintenance and firefighting. We look to best practices that result in well-running systems, a balanced workload for you and your team, and most importantly, a great customer experience.

The goal of operational excellence is to get new features and bug fixes into customers' hands quickly and reliably. Organisations that invest in operational excellence consistently delight customers while building new features, making changes, and dealing with failures. Along the way, operational excellence drives towards continuous integration and continuous delivery (CI/CD) by helping developers achieve high quality results consistently.

Design principles

- **Perform operations as code**: In the cloud, you can apply the same engineering discipline that you use for application code to your entire environment. You can define your entire workload (applications, infrastructure, etc.) as code and update it with code. You can script your operations procedures and automate their process by launching them in response to events. By performing operations as code, you limit human error and create consistent responses to events.
- **Make frequent, small, reversible changes**: Design workloads that are scalable and loosely coupled to permit components to be updated regularly. Automated deployment techniques together with smaller, incremental changes reduces the blast radius and allows for faster reversal when failures occur. This increases confidence to deliver beneficial changes to your workload while maintaining quality and adapting quickly to changes in market conditions.
- **Refine operations procedures frequently**: As you evolve your workloads, evolve your operations appropriately. As you use operations procedures, look for opportunities to improve them. Hold regular reviews and validate that all procedures are effective and that teams are familiar with them. Where gaps are identified, update procedures accordingly. Communicate procedural

updates to all stakeholders and teams. Gamify your operations to share best practices and educate teams.

- **Anticipate failure**: Perform "pre-mortem" exercises to identify potential sources of failure so that they can be removed or mitigated. Test your failure scenarios and validate your understanding of their impact. Test your response procedures to ensure they are effective and that teams are familiar with their process. Set up regular game days to test workload and team responses to simulated events.
- **Learn from all operational failures**: Drive improvement through lessons learned from all operational events and failures. Share what is learned across teams and through the entire organisation.
- **Use managed services:** Reduce operational burden by using AWS managed services where possible. Build operational procedures around interactions with those services.
- **Implement observability for actionable insights:** Gain a comprehensive understanding of workload behavior, performance, reliability, cost, and health. Establish key performance indicators (KPIs) and leverage observability telemetry to make informed decisions and take prompt action when business outcomes are at risk. Proactively improve performance, reliability, and cost based on actionable observability data.

While operational excellence design principles are focused on digital workloads, their wider objective is to help an organisation improve its operational capability. The following specific questions are intended to identify areas of improvements to an organisation's operational practices with respect to Māori data.

Topics

MD_OPS 1: How do you incorporate Māori views into your technology governance and operations?

Where appropriate, the Māori world view (te ao Māori) and Māori knowledge and understanding (mātauranga Māori) could be incorporated into an organisation's way of operating its technology in support of Māori customers. Sometimes this cannot be done as easily as performing operations as code, instead it may require active and ongoing collaboration with Māori. Understanding Māori interests and considerations relating to data collection, processing, and storage, as well as how those interests are evolving, helps you make informed technology decisions both in the interim and the long-term.

- **MD_OPS01-BP01: Incorporate mātauranga Māori into your operational processes when it relates to Māori data.** There are several ways of doing this. Work directly with your Māori customer or accessible internal and external Māori advisers to help develop your organisation's understanding of te ao Māori.
- **MD_OPS01-BP02: Consider governance and accountability over Māori data.** Consider incorporating specific accountabilities for Māori data into pre-existing roles in your organisation, such as Chief Information Officer (CIO), Chief Information Security Officer, or Chief Data Officer (CDO). Formalising these accountabilities and responsibilities into position descriptions can prioritise Māori data considerations at the executive level.
- **MD_OPS01-BP03: Consider how you can incorporate international data management principles into your data governance framework**. Some examples of international data management principles include Findability, Accessibility, Interoperability, Reuse (FAIR), which were developed for scientific data management and stewardship.
- **MD_OPS01-BP04: Consider how to supplement your knowledge for example through using Māori cultural advisers at the appropriate time**. Build your network of Māori advisers externally to your organisation, which you can engage with on particular technology projects. AWS has a growing list of advisers in the Amazon Partner Network. Working with external advisers also helps you develop your internal capabilities.
- **MD_OPS01-BP05: Develop longer-term policies and procedures.** Consider how to gradually develop the internal competencies and capabilities in your organisation as you continue to work with Māori customers. This can help you set your organisation's or your customer's overall data policies and processes with respect to collecting, storing, processing, and handling Māori data. Raising awareness and understanding of this across your technology staff and suppliers can help inform design decisions for handling Māori data, but this may be a long-term process.

MD_OPS 2: How can you design data collection with your Māori customer(s) in mind?

Organisations collect and process data to support the delivery of products and services to customers, stakeholders, and citizens. Data collection occurs in many ways, including filling in a digital form on a website, sensors capturing environmental data like temperature or water flow rates, or a research team capturing data from participants in a university study. Organisations need to consider what data they are collecting, the purpose for collecting the data, and, in the case of personal information, to adhere to the New Zealand Privacy Act 2020. For additional considerations around the collection of personal information, see Using AWS in the Context of New Zealand Privacy Considerations.

- **MD_OPS02-BP01: Consider adopting a privacy by design approach by designing and implementing mechanisms and processes that simplify compliance with the New Zealand Privacy Act 2020.** When collecting and handling personal information, make sure you comply with all applicable laws, such as the New Zealand Privacy Act 2020. You can design continuous and informed consent mechanisms that provide clear information to users about what personal information is being collected and for what purpose. Consider incorporating mechanisms that make it easier for users to revoke consent and to request access to, or correction of, their personal information. Maintain consent management over how you capture, store, and preserve personal information.
- **MD_OPS02-BP02: Consider how you're communicating why you are collecting this data.** Make it clear to users why the data is being collected, how it is used, and how privacy is maintained on an ongoing basis. Customers can interact with your organisation multiple times, so communicating what data is collected, why it is collected, and how it is collected in the context of the interaction may help. Also consider when to communicate this. Some examples include:
  - At the time a user registers or signs up to an organisation (for example, they sign up to a new bank, a new medical centre, or subscribe to a music streaming service).
  - At the time a user applies for a service from your organisation (for example, they apply for a home loan, book a medical exam, request a quote for home improvements, or apply for a government entitlement like a student loan).
  - At the time a user interacts with your organisation (for example, they lodge a complaint with a local council, submit an insurance claim, or request a change to a home loan).

- **MD_OPS02-BP03: Consider important lineage and provenance of data that could be captured as additional data.** It may be helpful to capture and store lineage and provenance data, which could be included as metadata or a tag. This kind of additional data can provide additional context to the data, such as when it was collected, how it was collected and who was involved in the collection. Maintaining careful records of the data provenance can build trust in the integrity and authenticity of the data and from whom the data was derived. An example of this is a hapū's cultural archive which captures which whānau provided which cultural record. It could also apply to organisations conducting surveys of specific populations for the purpose of creating data sets or performing data analysis.

MD_OPS 3: How do you use or share Māori data back with Māori?

Organisations use data in many different ways to design and deliver products and services. They can use it to gain insight and understanding of their organisation, their industry, and the wider world around them. Organisations should assess any legal and ethical implications when making decisions relating to how data will be used and shared.

- **MD_OPS03-BP01: Collecting and separating Māori data appropriately**. Organisations can consider how to collect and separate data along different dimensions such as iwi and hapū or a Māori organisation. This may make data more relevant or useful to different groups or communities. This needs to be considered when designing your initial data collection plan to verify that the right data is collected and organised early. For example, a government agency may want to report on specific agency outcomes for Māori populations. If they captured an individual's hapū affiliations, they could report information at a hapū level. However, if they only captured an individual's record, then they can only report this at an individual level.
- **MD_OPS03-BP02: Consider how your organisation could share Māori data back to Māori**. Data that your organisation holds could be used to better understand Māori communities and realise individual and collective benefits. Consider how you could identify useful data, and design ways to make data more accessible. Open data initiatives are one approach, but also remember the importance of complying with the Privacy Act 2020.
- **MD_OPS03-BP03: Use tools to effectively and securely share data where there is a specific and appropriate purpose.** There are many approaches to sharing data. Consider tools to share data both publicly and privately, which fits the purpose for why that data is being shared and how it needs to be used. Share this data in culturally-appropriate ways, and avoid misappropriating that data or trying to create explanation around that data that isn't culturally-

sensitive or informed. Seek advice from your Māori advisers if you are unsure how to do this. This includes open data portals or exchanges such as the AWS Open Data Registry, AWS Data Exchange, or private data portals hosted by your organisation. Consider solutions that allow data to be shared using different formats such as flat files, application programming interfaces (APIs), or interactive dashboards and visualisations to meet the needs of data consumers.

- **MD_OPS03-BP04: Share data in ways that can be easily used by your various stakeholders**. Understand your Māori stakeholders, their interests in accessing and using the data you are sharing back with them, and how they use the data. This should drive your choice of format for that data. For example, graphs or visualisations published on your website can make data easy to find and understand, while data files such as comma-separated values files (CSV) or parquet are better for data analytics users. An API supports application-to-application integration. For example, a government agency may provide interactive graphs on their website so that anyone with a web browser can see graphs relating to key agency objectives. They may provide the data that sits behind the graphs in a CSV format so that people can download it and load it into a spreadsheet tool or analytics programme to build their own graphs or perform additional queries. They may also provide an open API so members of the public can retrieve the data and load into their own databases or analytics systems.
- **MD_OPS03-BP05: Consider how your organisation could use federated data access methods.** Organisations often need to access data from other organisations to support a business process. Federated data approaches can allow organisations to access data from other organisations without the need to replicate or copy the data into your own systems. Federated data access models require appropriate mechanisms for making your organisations data discoverable and for securely sharing data.
- **MD_OPS03-BP06: Define and implement appropriate authorisation mechanisms.** Design data access mechanisms that support both internal access and access for external third parties (for example, through federated data access or data sharing mechanism). Consider authorisation mechanisms including role-based access control (RBAC), attribute-based access control (ABAC), or policy-based access control (PBAC). The authorisation mechanism should provide ways to manage, grant, and revoke access and provide visibility into data access through auditing and logging. Establish appropriate governance processes to effectively manage the process for requesting, granting, and revoking access to data by verified, trusted, and approved external third parties.

- **MD_OPS03-BP07: Incorporate responsible and ethical use of machine learning (ML) and artificial intelligence (AI) as a core part of your governance framework and development lifecycle.** ML and AI have transformational potential. It is already widely used for tasks such as

transcription, translation, fraud detection, information security, search, and recommendation engines. At Amazon, we believe the design, development, and deployment of AI must respect the rule of law, human rights, and values of equity, privacy, and fairness. We are committed to developing fair and accurate AI services and providing customers with the tools and guidance needed to build applications responsibly. Developers and deployers of AI systems should ensure such systems are built based on principles of safety and responsibility by design. AWS builds AI with responsibility in mind at each stage of our comprehensive development process. Throughout design, development, deployment, and operations, we consider a range of factors including accuracy, fairness, appropriate usage, toxicity, security, safety, and privacy.

- **MD_OPS03-BP08: Leverage vendor tools to provide AI transparency.** For example, AWS AI Service Cards deliver a form of transparency documentation that provide customers with a single place to find information on the intended use cases and limitations, responsible AI design choices, and deployment and performance optimisation best practices for our AI services. Amazon SageMaker Clarify detects and measures potential bias using a variety of metrics so developers can address potential bias and explain model predictions. Amazon's Responsible Use of Machine Learning Guide highlights key best practices and tooling that AI developers and deployers can use to mitigate risks across the lifecycle of an AI system.
- **Other considerations:**
  - **Discuss AI with stakeholders.** Artificial intelligence including generative AI and machine learning is a rapidly evolving area. Discuss the benefits and risks with your stakeholders. Your stakeholders may be internal to your organisation, external customers, or members of the public.
  - **Use confidence levels.** It's important to understand that many systems generate predictions of a possible or likely answer, not the answer itself. Confidence levels, if available, should be considered when reviewing outputs provided by the system.
  - **Bring human review into your systems.**
  - **Use case evaluation and testing.** Testing should include not just the AI system itself but also the overall process it is a part of, including decisions or actions that might be taken based on system output.
  - **Continuous improvement and validation.** Monitoring for potential bias and accuracy, and for models performing as expected across different segments, is an important part of this process.
  - **Ongoing education.** AI is a constantly-evolving landscape, and new techniques, technologies, laws, and social norms continue to be developed and refined over time. It is essential that all parties involved with building and using AI systems stay educated on these issues and account for them in the design, deployment, and operation of their systems.

For further reading, refer to the AWS Machine Learning lens.

AWS continues to update this information and share additional guidance to customers on the use of AI/ML. Please reach out to the team at AWS for further updates.

Resources

The following resources are specific to operational excellence and can help you apply Māori data considerations.

- FAIR principles
- Implementing FAIR principles into Health and Life science data lakes
- AWS Machine Learning lens
- AWS Responsible use of Machine Learning guide
- Using AWS in the Context of New Zealand Privacy Considerations

Security

The Security Pillar helps you meet your business and regulatory requirements by following current AWS recommendations. It's intended for those in technology roles, such as chief technology officers (CTOs), chief information security officers (CSOs/CISOs), architects, chief information/privacy officers (CIOs/CPOs), developers, and operations team members. After reading this document, you can understand current AWS recommendations and strategies to use when designing cloud architectures with security in mind. You can find prescriptive guidance on implementation in the Security Pillar whitepaper.

The Security Pillar describes how to take advantage of cloud technologies to protect data, systems, and assets in a way that can improve your security posture. The Security Pillar whitepaper provides in-depth, best-practice guidance for architecting secure workloads on AWS.

Design principles

- **Implement a strong identity foundation.** Implement the principle of least privilege and enforce separation of duties with appropriate authorisation for each interaction with your AWS resources. Centralise identity management, and aim to eliminate reliance on long-term static credentials.
- **Maintain traceability.** Monitor, alert, and audit actions and changes to your environment in real time. Integrate log and metric collection with systems to automatically investigate and take action.
- **Apply security at all layers.** Apply a defence in depth approach with multiple security controls. Apply to all layers (for example, edge of network, virtual private cloud (VPC), load balancing, every instance and compute service, operating system, application, and code).
- **Automate security best practices.** Automated software-based security mechanisms improve your ability to securely scale more rapidly and cost-effectively. Create secure architectures, including the implementation of controls that are defined and managed as code in version-controlled templates.
- **Protect data in transit and at rest**. Classify your data into sensitivity levels and use mechanisms, such as encryption, tokenisation, and access control where appropriate.
- **Prepare for security events.** Prepare for an incident by having incident management and investigation policy and processes that align to your organisational requirements. Run incident response simulations and use tools with automation to increase your speed for detection, investigation, and recovery.

The preceding design principles aim to support an organisation to improve its security posture. The following specific questions are some additional security considerations from a Māori data perspective.

Topics

MD_SEC 1: How is Māori data protected?

Systems designed to capture, store, or process Māori data should follow the same best practice as any other cloud solution in that they should be designed, built, and operated with security in mind. The Security Pillar of the Well-Architected Framework provides in-depth, best practice guidance for architecting secure workloads. The Data protection best practice area in the Security Pillar provides best practices relating to data classification, protecting data at rest, and protecting data in transit. Data protection is just one aspect of securing your cloud architectures. Security should be applied at all layers through multiple controls using a defence-in-depth approach. In addition to the best practices contained in the Security Pillar, the following considerations may also apply:

- **MD_SEC01-BP01: Design storage systems to handle data with different Māori data classifications.** If the data is considered tapu and that is feedback you have received from your customers, then it may be appropriate to apply additional controls and procedures. These may be required to support specific tikanga related to the handling of certain data. You may choose to store certain data separately from other data with different security requirements for access and processing. It is possible, for example, to store data in separate virtual private clouds (VPC), separate databases, or separate object storage buckets. This separation of datasets means you can apply independent security controls, such as different access permissions, different logging and auditing levels, and different backup approaches.

MD_SEC 2: How do you design workload security for long-term safety?

Certain Māori data may need to be available for generations to come. Consult with Māori customers and advisers about what data retention policies they recommend according to the different types of data. These data retention policies can be revisited in the future too. Regardless of how long you are intending to store this data, all data needs to be properly secured for the protection of taonga (treasure) for generations to come.

Ransomware is a good example to consider. If you have one copy of your data and you are subject to a ransomware attack, you may not be able to recover your data. Consider how many backup copies may be required to protect yourself from this scenario. Design appropriate access controls to minimise the chance of accidental or malicious deletion or corruption of back-ups. While it may seem redundant, it's important to store backups across multiple different types of storage and in multiple different locations. With this strategy, there's always an available backup, no matter the circumstances. Where irreplaceable digital taonga is identified, it is important to consider offline replication in addition to the appropriate data protection and resilience controls.

- **MD_SEC02-BP01: Understand data protection options available through your provider to protect data at the level of control your customer wants.** Customers control how they configure their environments and secure their content, including whether they encrypt their content (at rest and in transit), and what other security features and tools they use and how they use them. AWS does not change customer configuration settings, as these settings are determined and controlled by the customer. AWS customers have the complete freedom to design their security architecture to meet their compliance needs. AWS provides the customer autonomy to decide when and how security measures are implemented in the cloud, in accordance with each customer's business needs. When choosing which option is best, you should understand the risks you are trying to mitigate, take into account both the benefits and costs of each solution, and choose a solution that meets your requirements. Choose a cloud provider that offers contractual restrictions on their access to your data and operational restrictions. AWS, for example, is one of those cloud providers who offers both.
- **MD_SEC02-BP02: Understand what encryption options are available to protect your data at rest and in transit.** Encryption of data at rest is a recommended best practice for protecting your data from unauthorised access. AWS provides several options for data encryption. One option is to have AWS create and manage encryption keys for you through the AWS Key Management Service. Many AWS services integrate with AWS KMS to enable encryption of your data. Another option is to create your own encryption keys within AWS KMS. This provides you with more

control over your keys. This includes control over the key material, the rotation policy, and the permissions that define who can use or manage the key. AWS KMS is designed so that no one, not even an AWS employee, can retrieve your plaintext KMS keys from the service.

- **MD_SEC02-BP03: Make informed decisions about where data is stored**. Māori users of your system may prefer their data to be stored in New Zealand. AWS allows customers to control where their data is stored and processed, and your content won't be replicated or moved outside of your chosen AWS Region except as agreed by you. For customers in Aotearoa New Zealand, the options for storing data within New Zealand include the Auckland AWS Local Zone, an AWS Outpost, or the upcoming AWS Auckland region. Every commercial AWS region is designed, built, and operated in the same way and incorporates the same levels of security. When choosing an AWS infrastructure for your workload, take into account the possible trade-offs that may exist. For example, an AWS Region has a larger selection of AWS services and higher resiliency than an AWS Outpost. However, an AWS Outpost may provide more flexibility as to the location where the infrastructure can be placed. There are also costs and budget considerations to take into account. Some other considerations related to the location of data include:
  - Do you need to make a distinction between where data is processed and where it is stored? For example, the data could be stored in a database in New Zealand but processed on an EC2 instance in another region (of your choice) as part of an analytics job. Alternatively, it could be captured using a web application running on servers in Sydney and then saved to a database located in New Zealand.
  - Do you need to duplicate data across locations to meet your customer requirements? For example, a data archiving solution could send backups to another AWS Region for resiliency and security reasons. An application like a digital archive solution could make use of Amazon CloudFront for content distribution to help reduce latency for end users when accessing the content. This would require copies of data to be stored at CloudFront edge locations, while the primary data is stored in the origin storage service such as Amazon S3 or a database.

MD_SEC 3: How can you identify and classify Māori data?

Organisations may wish to understand what Māori data they hold and have a method to classify Māori data to protect it with security controls and practices. Once the data has been classified and appropriate metadata is captured, you can govern and use that data in appropriate ways.

- **MD_SEC03-BP01: Develop an understanding of what Māori data is**. Create an easy-to-understand definition of Māori data for your organisation. Having a definition of Māori data can help determine what additional considerations are relevant to your organisation or application.

Piloting this and testing this with your customers facilitates a mutually-agreed upon definition, which is implementable for your organisation.

- **MD_SEC03-BP02: Incorporate Māori data classification into your data governance framework**. If you already have a data classification framework within your organisation, you may wish to expand this to include a Māori data classification approach. A classification framework can help you determine what is Māori data, outline where and how the classification should be recorded, and define the security and access controls that are required for that data classification. For example, what additional data access controls may be required if the data is classified as tapu or noa?
- **MD_SEC03-BP03: Record Māori data classifications as metadata and make it easily discoverable**. Use approaches such as metadata tagging and data cataloguing to store and manage your classification metadata. Tools such as business data catalogues should make it easier for your organisation to search for and discover datasets that contain Māori data in your organisation. In addition, you can add tags which record the purpose for which consent of that data was given. This may allow for easier review of consent and data access policies and up-to-date consent practices.
- **MD_SEC03-BP04: Leverage technologies and techniques to help identify and classify existing Māori data**. Your organisation may already capture and store Māori data. Consider using available tools and techniques to review and classify your existing data. Once identified and classified, update your metadata and business data catalogues. This may involve assessing data across databases, object stores, file servers, document management systems, communication systems (like email), and analytics platforms. For example, search your AWS Glue Data Catalog table columns for terms like *iwi* or *hapū* or Māori organisation.

MD_SEC 4: How do you maintain privacy of personal Māori data?

It is important that personal data is collected and processed lawfully, fairly, and transparently in relation to a person. When designing, building, and operating workloads that may capture, store, and process Māori personal data, privacy should be taken into account throughout the entire process. The application of privacy principles should align with your organisation's privacy framework and be guided by applicable privacy regulation such as the New Zealand Privacy Act 2020. For further information on how you are applying AWS services in conjunction with the New Zealand Privacy Act 2020, see Using AWS in the Context of the New Zealand Privacy Considerations.

- **MD_SEC04-BP01: Use tools and techniques to adequately de-identify data.** This helps protect individual's privacy when producing data sets that may be shared or published. There are many techniques within data and analytics domains to help de-identify data. These include obfuscation (obscuring sensitive data), tokenisation (where a sensitive piece of data is replaced by a non-sensitive token where the token can map back to the original data), and anonymisation (such as removing sensitive data completely).
- **MD_SEC04-BP02: Honour the consent you've received when using data for internal analytics**. If your organisation asked for consent when collecting data, that data should be used only for the purposes that you have received consent for. If your organisation asked for consent and did not receive it, exclude this data from analytics and AI/ML uses.
- **MD_SEC04-BP03: Honour the consent you've received when sharing data**. Only share data with third parties if you have obtained requisite consent in accordance with applicable laws such as the New Zealand Privacy Act. Consider creating mechanisms to exclude that data from any data sharing processes if you have not obtained the required consents.

Resources

The following resources are specific to security and can help you apply Māori data considerations.

- AWS Well Architect Framework – Security Pillar
- Data Residency Whitepaper
- AWS Compliance site
- Responsible Use of Machine Learning guide
- AWS Data Analytics Lens
- AWS Government Lens – Privacy by Design
- The Security Design of the AWS Nitro System: No AWS Operator Access
- AWS Nitro System gets independent affirmation of its confidential compute capabilities

Reliability

The Reliability Pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to. This includes the ability to operate and test the workload through its total lifecycle. This document provides in-depth, best practice guidance for implementing reliable workloads on AWS. If relevant, you can find prescriptive guidance on implementation in the Reliability Pillar whitepaper. This pillar encompasses the ability of a workload to perform its intended function correctly and consistently when it's expected to.

Design principles

- **Automatically recover from failure:** By monitoring a workload for key performance indicators (KPIs), you can run automation when a threshold is breached. These KPIs should be a measure of business value, not of the technical aspects of the operation of the service. This allows for automatic notification and tracking of failures, and for automated recovery processes that work around or repair the failure. With more sophisticated automation, it's possible to anticipate and remediate failures before they occur.
- **Test recovery procedures:** In an on-premises environment, testing is often conducted to prove that the workload works in a particular scenario. Testing is not typically used to validate recovery strategies. In the cloud, you can test how your workload fails, and you can validate your recovery procedures. You can use automation to simulate different failures or to recreate scenarios that led to failures before. This approach exposes failure pathways that you can test and fix *before* a real failure scenario occurs, thus reducing risk.
- **Scale horizontally to increase aggregate workload availability:** Replace one large resource with multiple small resources to reduce the impact of a single failure on the overall workload. Distribute requests across multiple, smaller resources to ensure that they don't share a common point of failure.
- **Stop guessing capacity:** A common cause of failure in on-premises workloads is resource saturation, when the demands placed on a workload exceed the capacity of that workload (this is often the objective of denial of service attacks). In the cloud, you can monitor demand and workload utilisation, and automate the addition or removal of resources to maintain the optimal level to satisfy demand without over- or under-provisioning. There are still limits, but some quotas can be controlled and others can be managed.
- **Manage change through automation**: Changes to your infrastructure should be made using automation. The changes that need to be managed include changes to the automation, which then can be tracked and reviewed.

From a Māori data perspective, there are other priorities and additional views of why reliability is important and how to go about improving reliability. The following specific questions and good practices complement best practices in the Reliability Pillar.

Topics

MD_REL 1 How do you safely retain data for future generations?

Māori data is often considered taonga, and it is critical that this data is available for future generations. Whakapapa (genealogy) or mātauranga Māori (Māori knowledge) are examples where long-term retention and protection is important.

- **MD_REL01-BP01: Design storage systems with multi-generational durability in mind**. Multi-generational durability focuses on the reliability of having access to that data across generations. This applies to the data that you are storing, but also associated metadata that provides context for that data. The Well-Architected Reliability Pillar provides guidance on best practices for architecting resilient workloads, backing up data, and planning for disasters. By understanding what kinds of data you are storing and the possible need for the long-term preservation of that data, you can choose appropriate architectural patterns and services such as Amazon S3, which creates six copies of your data and is designed to provide 99.999999999% (11 nines) of durability. In practice, 11 nines of durability means that if you stored ten million objects, you might expect to lose a single object every 10,000 years.
- **MD_REL01-BP02: Consider how your organisational archive and retention processes apply to Māori data.** Many organisations have data retention and archiving processes that govern how long data is retained and how and where it is stored.
- **MD_REL01-BP03: Configure your AWS account for long term continuity**. Within each of your AWS accounts, it is important to have accurate and up-to-date contact details, payment/credit card information, and multi-factor authentication (MFA) for users. Keeping your contact information up-to-date helps ensure that you receive important notifications from AWS on topics like security, billing, and operations. It is a best practice to use an email distribution list, rather than depending on an individual's email address. This can help avoid scenarios

where important notifications from AWS are missed if an individual is on leave or has left the organisation.

- **MD_REL01-BP04: Implement backup mechanisms**. The Reliability Pillar provides guidance on designing, implementing, and operating a backup solution for your data and applications. The nature of the Māori data being captured, processed, or stored influences the design of the backup solution. For example, an iwi register application or a digital archive application might call for multiple copies of backups to be stored in different locations, with different access controls due to the value of those datasets. While it may seem redundant, it's important to store backups across multiple different types of storage and in multiple different locations. This helps ensure there's always an available backup, no matter the circumstances. Where irreplaceable digital taonga is identified, it is important to consider offline replication in addition to the appropriate security and reliability controls.

Resources

The following resources are specific to reliability and can help you apply Māori data considerations.

- AWS Well Architected Framework – Reliability Pillar
- AWS Data Analytics lens

Performance efficiency

The Performance Efficiency Pillar addresses best practices for managing production environments. This document does not cover the design and management of non-production environments and processes, such as continuous integration or delivery. The Well-Architected Performance Efficiency Pillar provides an overview of design principles, best practices, and questions. If relevant, you can find guidance on implementation in the Performance Efficiency Pillar whitepaper.

The Performance Efficiency Pillar focuses on the efficient use of computing resources to meet requirements, and how to maintain efficiency as demand changes and technologies evolve. Following these design principles can help you achieve and maintain efficient workloads in the cloud.

Design principles

- **Democratise advanced technologies:** Make advanced technology implementation easier for your team by delegating certain tasks like the undifferentiated work to your cloud vendor. Rather than asking your IT team to learn about hosting and running a new technology, consider consuming the technology as a service. For example, NoSQL databases, media transcoding, and machine learning are all technologies that require specialised expertise. Learn and develop these expertise with AWS Skill Builder. In the cloud, these technologies become services that your team can consume, allowing your team to focus on product development rather than resource provisioning and management.
- **Go global in minutes:** Deploying your workload in multiple AWS Regions around the world allows you to provide lower latency and a better experience for your customers at minimal cost.
- **Use serverless architectures:** Serverless architectures remove the need for you to run and maintain physical servers for traditional compute activities. For example, serverless storage services can act as static websites (removing the need for web servers) and event services can host code. This removes the operational burden of managing physical servers, and can lower transactional costs because managed services operate at cloud scale.
- **Experiment more often:** With virtual and automatable resources, you can quickly carry out comparative testing using different types of instances, storage, or configurations.
- **Consider mechanical sympathy:** Use the technology approach that aligns best with your goals. For example, consider data access patterns when you select database or storage approaches.

Cost optimisation

The Cost Optimization Pillar provides guidance on a cost-optimised workload. A cost-optimised workload fully uses all resources, achieves an outcome at the lowest possible price point, and meets your functional requirements. This document provides in-depth guidance for building capability within your organisation, designing your workload, selecting your services, configuring and operating the services, and applying cost optimisation techniques. You can find prescriptive guidance on implementation in the Cost Optimization Pillar whitepaper.

Cost optimisation is a continual process of refinement and improvement over the span of a workload's lifecycle. The practices in this document help you build and operate cost-aware workloads that achieve business outcomes while minimising costs and allowing your organisation to maximise its return on investment.

Design principles

- **Implement cloud financial management:** To achieve financial success and accelerate business value realisation in the cloud, you must invest in Cloud Financial Management. Your organisation must dedicate the necessary time and resources for building capability in this new domain of technology and usage management. Similar to your Security or Operations capability, you need to build capability through knowledge building, programs, resources, and processes to help you become a cost efficient organisation.
- **Adopt a consumption model:** Pay only for the computing resources you consume, and increase or decrease usage depending on business requirements. For example, development and test environments are typically only used for eight hours a day during the work week. You can stop these resources when they're not in use for a potential cost savings of 75% (40 hours versus 168 hours).
- **Measure overall efficiency:** Measure the business output of the workload and the costs associated with delivery. Use this data to understand the gains you make from increasing output, increasing functionality, and reducing cost.
- **Stop spending money on undifferentiated heavy lifting:** AWS does the heavy lifting of data center operations like racking, stacking, and powering servers. It also removes the operational burden of managing operating systems and applications with managed services. This allows you to focus on your customers and business projects rather than on IT infrastructure.
- **Analyse and attribute expenditure:** The cloud makes it easier to accurately identify the cost and usage of workloads, which then allows transparent attribution of IT costs to revenue streams and

individual workload owners. This helps measure return on investment (ROI) and gives workload owners an opportunity to optimise their resources and reduce costs.

Topics

MD_CO 1 Understand costs associated with infrastructure options

- **MD_CO01-BP01: How are you presenting the costs/benefit tradeoffs for infrastructure options?** Clearly present the cost to benefit trade-offs when looking at all infrastructure options. When designing technology solutions, organisations need to balance delivering functional and non-functional requirements with the cost associated with doing so. It is important that the cost to benefit ratio is assessed and understood. The assessment should take into account the strategic drivers for the organisation. Within the public sector, for example, value-for-money in terms of social, environmental, and public benefits is important. For some Māori customers, security, resilience, and sustainability may be key considerations.
  - An example may be an organisation wanting to store data close to the source for latency reasons. An AWS Outpost can help meet those requirements. An AWS Outpost is a managed service that extends AWS infrastructure, services, APIs, and tools to customer premises. However, there are several considerations when using an AWS Outpost. An Outpost has a limited set of AWS services available compared to the 200+ services available in an AWS Region. This consideration may impact how you design your architecture. An AWS Outpost is a physical device that needs to be built, shipped, and installed within a datacentre that meets certain physical requirements. This can be done with a third-party data centre, but there is a cost associated with renting space.

Sustainability

The Sustainability Pillar provides design principles, operational guidance, best-practices, potential trade-offs, and improvement plans you can use to meet sustainability targets for your AWS workloads. If relevant, you can find prescriptive guidance on implementation in the Sustainability Pillar whitepaper.

Apply these design principles when architecting your cloud workloads to maximise sustainability and minimise impact.

Design principles

- **Understand your impact:** Measure the impact of your cloud workload and model the future impact of your workload. Include all sources of impact, including impacts resulting from customer use of your products, and impacts resulting from their eventual decommissioning and retirement. Compare the productive output with the total impact of your cloud workloads by reviewing the resources and emissions required per unit of work. Use this data to establish key performance indicators (KPIs), evaluate ways to improve productivity while reducing impact, and estimate the impact of proposed changes over time.
- **Establish sustainability goals:** For each cloud workload, establish long-term sustainability goals such as reducing the compute and storage resources required per transaction. Model the return on investment of sustainability improvements for existing workloads, and give owners the resources they need to invest in sustainability goals. Plan for growth, and architect your workloads so that growth results in reduced impact intensity measured against an appropriate unit, such as per user or per transaction. Goals help you support the wider sustainability goals of your business or organisation, identify regressions, and prioritise areas of potential improvement.
- **Maximise utilisation:** Right-size workloads and implement efficient design to ensure high utilisation and maximise the energy efficiency of the underlying hardware. Two hosts running at 30% utilisation are less efficient than one host running at 60% due to baseline power consumption per host. At the same time, eliminate or minimise idle resources, processing, and storage to reduce the total energy required to power your workload.
- **Anticipate and adopt new, more efficient hardware and software offerings:** Support the upstream improvements your partners and suppliers make to help you reduce the impact of your cloud workloads. Continually monitor and evaluate new, more efficient hardware and software offerings. Design for flexibility to allow for the rapid adoption of new efficient technologies.

- **Use managed services:** Sharing services across a broad customer base helps maximise resource utilisation, which reduces the amount of infrastructure needed to support cloud workloads. For example, customers can share the impact of common data centre components like power and networking by migrating workloads to the AWS Cloud and adopting managed services, such as AWS Fargate for serverless containers, where AWS operates at scale and is responsible for their efficient operation. Use managed services that can help minimise your impact, such as automatically moving infrequently accessed data to cold storage with Amazon S3 Lifecycle configurations or Amazon EC2 Auto Scaling to adjust capacity to meet demand.
- **Reduce the downstream impact of your cloud workloads:** Reduce the amount of energy or resources required to use your services. Reduce or eliminate the need for customers to upgrade their devices to use your services. Test using device farms to understand expected impact and test with customers to understand the actual impact from using your services.

Māori have deep connections to the land and the natural world. The land can act as a foundation for whakapapa, and as such can be an integral part of a person or community's identity. The natural world also provides resources, food, and shelter for Māori. This connection with the natural world means that it needs to be cherished and protected. Having a deeper understanding of this can inform your organisation's approach to reducing or minimising the impact of your technology decisions on the environment. Specifically, aim to incorporate the design principles, guidance, and best practices in the Sustainability Pillar into your organisation's approach to sustainability. The following specific questions and good practices could be considered along with the best practices in the Well-Architected Sustainability Pillar whitepaper.

Topics

MD_SUS 1 How do you design and operate systems to minimise potential impacts on the environment?

- **MD_SUS01-BP01: Develop an understanding of the unique relationship between Māori and the land**. As tangata whenua (people of the land), Māori have deep connections to the land and the natural world. Whenua is seen as taonga that must be protected for future generations. It is important to understand the possible impacts your technology decisions may have on the

environment. Understand your Māori customers in the context of their natural environment, their priorities, and their interests, and incorporate their perspectives into your technology when considering opportunities and environmental impact.

- **MD_SUS01-BP02: Consider how your solution could positively impact the environment.** Technology is often used to improve processes by making them more efficient or effective. When designing products, consider how you can incorporate features that could have direct or indirect positive impacts on the environment. For example, if you are designing a farm management system that tracks nitrate usage on paddocks to maximise pasture growth, consider a feature that highlights potential risk of nitrates making their way into farm waterways. Similarly, you can reduce landfill waste by using machine learning in your manufacturing process to help determine the optimal use of resources.
- **MD_SUS01-BP03: Consider where you might store and run your workload to make the most of the sustainability of the cloud.** A study by 451 Research shows that "moving IT workloads from on-premises data centres to the cloud would improve energy efficiency and reduce associated carbon emissions by nearly 80% on average" (among 515 surveyed enterprises across Japan, South Korea, Singapore, Australia, and India). 451 Research also estimates that emissions savings for organisations moving from on-premises data centers to cloud could increase to 93% if cloud data centers in APAC were powered by 100% renewable energy. This advantage is attributable to the combination of more energy-efficient servers and much higher server utilisation. As New Zealand customers move their workloads from enterprise data centres to the AWS Cloud, the carbon footprint of these workloads is reduced due to lower energy consumption. Furthermore, the AWS Asia Pacific (Auckland) Region is planned to be powered by 100% renewable energy. This is part of Amazon's Climate Pledge whereby Amazon is on path towards powering its infrastructure operations by 100 percent renewable energy by 2025.
- **MD_SUS01-BP04: Measure and report on carbon footprint.** AWS offers customers free use of the Customer Carbon Footprint Tool (CCFT). The CCFT provides the following features:

  1. Simple data visualisations to report on the emissions from your AWS usage following Greenhouse Gas (GHG) Protocol standards;
  2. Analysis of the changes in your emissions over time as you migrate workloads to AWS;
  3. Helps you re-architect applications, or deprecate unused resources; and
  4. Forecasts how your emissions change across your sustainability journey as Amazon progresses toward powering operations with 100% renewable energy.

You can also talk to your account executive about the range of tools available on the AWS Marketplace, which can support tracking and reporting of your organisation's sustainability data. Such data provides necessary insight to achieve your organisation's sustainability goals.

Resources

The following resources are specific to sustainability and can help you apply Māori data considerations.

- AWS Well Architected – Sustainability Pillar
- AWS customer carbon footprint calculator
- Sustainability at Amazon
- Sustainability in the Cloud
- Water Stewardship

Scenario: Digital archive solution

The following example explains a digital archive solution scenario, its requirements, relevant Māori data considerations, and how to apply the AWS Well-Architected Pillars.

A local iwi organisation has engaged a software development company to develop and present a high-level solution for a digital archive solution. The iwi is digitising their current archive of documents and images, as well as creating recordings of their oral history. They are also generating new digital content such as videos and photos that they want to preserve. The archive is difficult to manage at present, as it is stored in different places, including personal cell phones, online storage services, and external hard drives. This makes it difficult to find content, and there are concerns regarding the potential for data loss.

The following are requirements for the solution.

1. We want to be able to have a place where we can upload digital taonga like videos, documents (including emails and digital scans of paper documents), and photos.
2. We want our information managers to be able to easily organise, manage, and access content to support iwi business and information requests from members, as well as other iwi or hapū organisations.
3. We want to be able to restrict access to content to different groups of users.
4. We want to allow members to be able to easily contribute content that they may have, such as photos or videos from events or gatherings.
5. We want to be able to choose certain items to publish onto our website so both members and the public can discover and enjoy them appropriately.
6. The items that the digital archive system hold are extremely valuable to the iwi, so it must be secure from things like hacking and accidental deletion or corruption.

Solution concept

The development team has taken the initial set of requirements are put together a high-level solution concept. They have also created a high-level architecture that includes recommended AWS services.

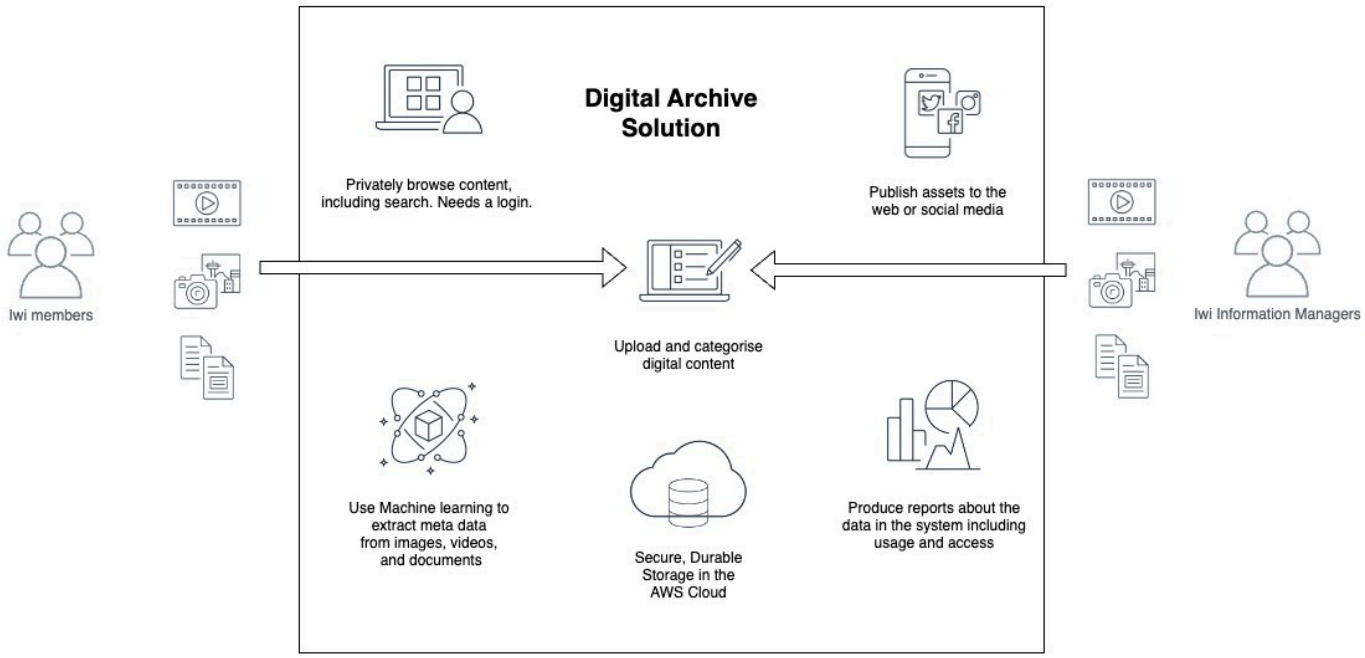

High-level architecture

The following diagram provides a high-level architecture for the digital archive solution.

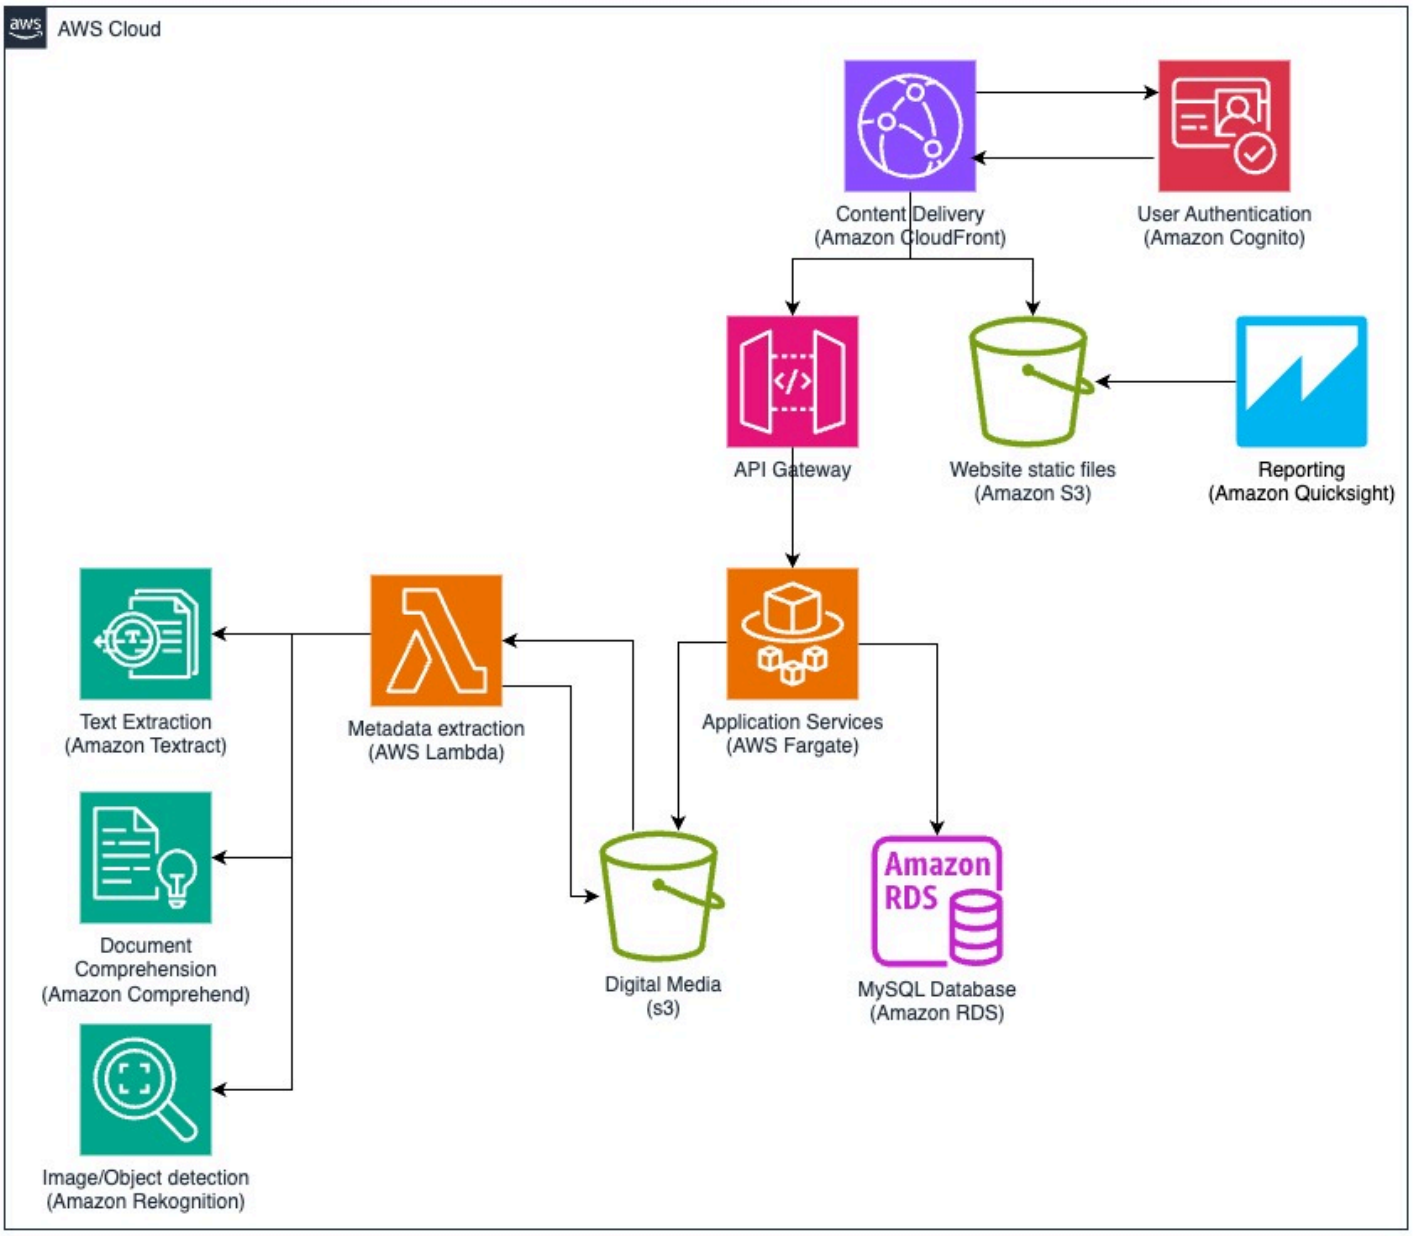

1. The front-end web application is a javascript based web application deployed to Amazon S3 and accessed using Amazon CloudFront. Amazon CloudFront is integrated with Amazon Web Application Firewall (WAF) to provide protection from layer 7 style attacks. Amazon CloudFront provides distributed denial of service (DDoS) protection through AWS Shield.
2. The web application uses Amazon Cognito for user authentication. External users can use existing social logins such as Facebook or Google or set up a new identity within the application. Amazon Cognito can be integrated with the iwi's existing identity provider using OAuth or SAML so that administrators and information managers can use existing identities if they exist. If one

does not exist, iwi users can have a new digital log in or digital identity record set up in Amazon Cognito.

3. The web application interacts with the back-end services using APIs exposed through Amazon API Gateway.
4. Application services and APIs are containerised using Amazon Elastic Container Service (ECS) and deployed to AWS Fargate. The APIs provide access to data, content, and support features such as searching, retrieving content from the content store, retrieving metadata from the database, content uploading, and user access management.
5. Uploaded items such as videos, images, and documents are stored in Amazon S3. This provides secure, durable, and cost-effective storage for digital content. Item metadata (such as source, name, description, date, location, and keywords) and other data such as user profiles, access roles and permissions, system usage, and auditing data are stored in a MySQL database. This database is deployed onto Amazon RDS for MySQL in a multi-AZ configuration to provide additional resilience.
6. AWS AI services are used to automatically extract useful metadata from uploaded content. This metadata can be used when searching for items in the archive. Extracted metadata is stored in the MySQL database. The type of metadata extraction depends on the type of item uploaded, but could include the following.
    - **Text extraction:** Text is extracted from documents using Amazon Textract.
    - **Document comprehension:** Key entities, such as people, organisations, or places, contained in documents is extracted using Amazon Comprehend. There is also the option to help classify documents or items.
    - **Object detection:** Amazon Rekognition is used to detect objects with images and videos, which can then be stored as meta-data.

Application of the Māori data guidance

This section outlines how the Māori data guidance could be applied to the digital archive solution scenario and the proposed solution architecture. It is written from the perspective of the software development company.

Operational excellence

How do you incorporate Māori views into your technology governance and operations?
This section focuses on developing general knowledge of te ao Māori within your organisation,

especially as it relates to how your organisation works with Māori as customers. In this scenario, consider the following:

- What level of Māori cultural capability does you organisation have to respond to this scenario? Is it enough to help you understand some of the needs and requirements that the iwi customer may have? If not, in the short term you may consider engaging a Māori expert to help your organisation work with this specific customer. In the longer term, you can consider developing this knowledge in your organisation.
- If your offering incorporates aspects of ongoing operations or support for the application, what can you do to incorporate the needs of your Māori customer? This may include the need to develop processes that provide the flexibility to support specific tikanga that your Māori customer may wish to incorporate into operational and support processes. For example, in the process of responding to a support call, staff may need to access part of the system that stores tapu (sensitive) data. The customer may wish that a specific protocol be followed by staff when accessing this tapu data.

How can you design data collection with your Māori customers in mind? This section focuses on how data is captured or collected. In this scenario, consider the following:

- The system allows external people to register for an account to use the system. Consider what data is required to allow a user to sign up and use the system. In this scenario, it may be appropriate to allow a user to indicate what hapū they associate with, as this could be used to help present relevant content to that user. Alternatively, rather than collecting data about a user, you could allow the user to specify their interests. These interests could then be used to personalise the content suggested to the user. At the time of registration, any relevant personal information collection notices should be clearly presented to the user in easy-to-understand language.
- The system allows users to make submissions for inclusion in the archive. For example, an authorised user could submit photos or images from an event they attended or submit historical documents or images that their whānau has collected over time. Given the solution is a digital archive, there should be requirements relating to the capture of metadata, like capturing where the data came from or what kind of rights the owner has given the holder of that information. From a Māori data perspective, the owner may have specific requirements for accessing or handling the data. This may require certain access and use restrictions to be put in place.

How do you use or share Māori data back with Māori? This section focuses on how Māori data is used or shared. Use and sharing can be from the perspective of a Māori organisation using the data, as is the case in this scenario, but it could also apply to third parties who collect, store, or generate Māori data, such as a medical centre, a government agency, or a non-profit delivering community services. The archive system in this scenario is designed to capture and store digital items by the iwi organisation. Some of the considerations in this section are from the perspective of the iwi and how they use or share the data captured and stored in the archives system.

- One of the key objectives for the archive solution is to make data accessible, and the features of the solution reflect this. In this scenario, the consideration of how Māori data could be shared back with Māori might mean how information from the iwi archive could be shared with other Māori organisations. For example, an individual hapū may have their own archives which could complement data from the broader iwi archive. Alternatively, data about who has accessed and seen the iwi archive data may be useful for hapū to understand the level of interest or engagement with the content they have shared.
- Consideration should be given to who might want to get data out of the archive system, as well as the most appropriate way retrieve that data. For example, the solution may allow individual users to do a search for content and download the specific item with its metadata. But what happens in a scenario where a third party may want to get multiple items? For example, if an individual hapū also has an archive system, how could they integrate their system with the iwi archive system? An API could facilitate programmatic access, which may make it easier to integrate and retrieve large amounts of content.
- The proposed architecture makes use of AWS AI services to perform tasks such as text extraction, document comprehension, transcription of video and audio, and object recognition in images and videos. The hapū can decide how the service operates in respect to the inputs and outputs. Does the service retain the inputs (like documents and images) for any purpose, and if so, who has access to the inputs? Are inputs used to develop and improve the service, and if so, is there a way to opt out of that transfer? For example, Amazon Transcribe, Amazon Comprehend, Amazon Textract, Amazon Rekognition, and Amazon Translate all allow customers to opt-out of the transfers of customer data to develop and improve services. For more detail, see Privacy Features of AWS Services.
- Because the architecture makes use of AWS AI services, work with the iwi customer to verify that they understand how the proposed AI services work, what function they serve, and what the benefits of the AI services are. You should seek guidance from them on the suitability of using such tools to perform the specific functions and where tikanga may need to be applied. For example, one consideration may be separation between the living and dead. Historical content

often relates to those who have passed, so there may be a desire to process items separately or even exclude certain items from being processed by AI services. Another consideration is the accuracy of AI services when analysing te reo Māori in written or spoken form. If the AI services have not been trained on te reo Māori, the transcriptions produced, entities identified, or classifications determined may be inaccurate or incomplete, which can reduce the usefulness of that the data. The system architecture or features need to incorporate the outcomes of these discussions. For example, there may be logic that checks for a specific tag on an item and uses that to determine if an item should be sent to AI services for text extraction or object detection.

Security

How is Māori data protected? This section focuses on specific security considerations from a Māori data perspective.

- The archive system holds a range of data, some of which may be considered tapu, and therefore data restrictions may need to be in place. There may be a need to store tapu data separately from other data. Guidance is required from the iwi customer on what separation looks like in a digital system. The preceding high-level architecture diagram shows that Amazon S3 and Amazon RDS are used as data stores. Information that is classified as tapu may need to be stored in a separate S3 bucket. The application would require logic to determine which bucket to save items into and provide functions that move the data between buckets if the data classification changed. The Amazon RDS database stores metadata about the item. This may include data about people, places, and events. If some of the metadata is classified as tapu, guidance should be sought on whether the data needs to be stored in a different database table or possibly a separate database. This would then be balanced with potential system complexity and cost.
- The classification of the data may also require additional access and security controls. Restricting access could be achieved through role-based or attribute-based access controls. System administrators could control access to more sensitive data through the granting of permissions to a user or role. Typically, audit logging provides traceability of who has accessed items in the archive. This can be used to validate the security access controls are working.
- Protecting data for long-term safety can be achieved by incorporating AWS Well-Architected security and resiliency best practices. From a security perspective, this includes understanding the threats that your application and organisation face. Identify mitigations that can be implemented as security controls. Given the proposed architecture for the archive solution, a potential threat to the long-term safety of the digital content stored in Amazon S3 is a ransomware event. Once a potential event is identified, determine steps that can be taken to

help protect your application, detect if this kind of event occurs, and respond to, and recover from such an event. For more detail, see The anatomy of ransomware event targeting data residing in Amazon S3.

How can you identify and classify Māori data? This section focuses on understanding what Māori data is in the context of your organisation and having a method to classify data as Māori data. This can then guide your architecture when capturing, processing, and storing that data. In this scenario, consider the following:

- It's clear that the archives system contains Māori data, considering that the customer is an iwi organisation and the system stores and processes data about their history, knowledge, and people.
- The high-level requirements include being able to control access to data for different users. This indicates that there are different types of data stored within the archive. Discuss how data is classified and how that classification is recorded in the system with the iwi customer. Is there an expectation that the system automatically determines the classification based on the document content or metadata? Should a user manually classify the document? Once classified, how do you record the classification as a piece of metadata, and link this to the digital item?
- The other considerations in this section are mainly for organisations that capture and process Māori data as part of delivering their products and services. They do not apply to this scenario.

How do you maintain the privacy of personal Māori data? This section focuses on maintaining privacy of personal data. In this scenario, consider the following:

- The application is likely capturing personal data when a new user registers with the application. Consideration needs to be given to how much data is collected, for what purpose, and how this is communicated to new users as well as how ongoing consent gets managed so that a user has the option to revoke access to that personal information and that all collection and management of personal information is in accordance with New Zealand privacy laws.

Reliability

How do you safely retain data for future generations? This section focuses on understanding that Māori data often needs to be protected and resilient so it can be accessed by future generations. In this scenario, consider the following:

- The archive holds extremely valuable taonga for the iwi organisation. Given its importance, the architecture needs to ensure that the data is resilient over time. The digital content is stored in Amazon S3, which provide high levels of durability (11 nines). This provides a level of protection from data loss caused by service events. Amazon S3 features like versioning can be used to protect data from deletion or corruption events. For more information on versioning, see Using Versioning in Amazon S3 buckets. Regular backups using Amazon S3 replication or AWS Backup can provide another layer of protection by creating additional copies of the objects stored in the archive.
- Archive and preservation systems often have multiple copies of content. There may be a preservation master and one or more copies that are used for general access. In some cases, the general access copies may be modified by lowering the resolution of video to be more easily consumed on a variety of devices or converting documents from one format to another to make it easier to consume on a range of different devices. Supporting multiple copies has cost implications in terms of storage. It also requires application features to perform functions such as file copies, file conversions, and image or video resampling.
- The proposed architecture uses an Amazon RDS database to store system data, including data about users, system usage, and metadata about items in the archive. To protect this database, the native database backup feature can be used to create regular backups of the database. Operational processes need to be established to verify that backups are occurring as expected and test the restoration process periodically.

Cost optimisation

This section focuses on understanding the cost considerations when designing a solution. In this scenario, consider the following:

Clearly present the cost to benefit trade-offs when looking at all infrastructure options. There may be a desire to have data located close to the iwi organisation. At the time of writing (June 2024), the nearest AWS Regions to New Zealand are Sydney and Melbourne, Australia. The Auckland Local Zone is available and is parented to the AWS Sydney Region. AWS Outposts is a fully managed service that extends AWS infrastructure, services, APIs, and tools to customer and data centre provider premises. An AWS Outpost could be deployed into a data centre close to the iwi organisation. Cost components include the AWS Outpost and the cost of hosting the AWS Outpost in a third-party data centre that meets the minimum requirements. The proposed architecture also uses AWS AI services, including Amazon Textract, Amazon Comprehend, and Amazon Rekognition. These services currently are only available in an AWS Region. The solution

would therefore need to consider the network connectivity and bandwidth requirements from the AWS Outpost to the Region, which may impact the overall solution cost.

Sustainability

How do you design and operate systems to minimise potential impacts on the environment? This section focuses on considering the impacts of technology on the environment.

- Work with your customer to identify if they have specific sustainability goals, and identify what metrics are being used to measure attainment of those goals. Determine how you might produce data from the digital archive solution that can feed into the measurement of those metrics.
- You may prompt the iwi to consider how the iwi can reduce carbon emissions by using AWS instead of alternatives like on-premise servers. You may wish to discuss with the iwi's kaitiaki board the pros of using a monthly report from the AWS Customer Carbon Footprint Tool to monitor carbon emissions and set a 12 month goal to reduce carbon emissions associated with their use of AWS through optimisations.

Conclusion

The goal of the Māori Data Lens is to provide guidance to organisations designing and operating solutions on AWS on how to apply a Māori data perspective to their workloads. The guidance was developed with Māori data experts and cultural advisors. It is designed to be applied in conjunction with the core Well-Architected pillars, principles, and best practices. Organisations that are designing solutions for Māori or solutions that capture, process, and store data about Māori can use the guidance to understand some of the unique considerations and perspectives that can be reflected in your solutions and organisation. The guidance is not intended to be an authoritative, prescriptive checklist. Instead, we expect organisations to partner with Māori advisors to develop their organisational understanding and competency. The Māori Data Lens is one tool that can be used to incorporate that understanding into your solution design processes. The application of the guidance will take time, reflection, and a process you can apply to your Māori customers, organisation, industry, and workload.

Finally, our call to action is the following:

- Determine what Māori data you have
- Design your cloud and use of Māori data with mechanisms, like those suggested in this document
- Support your business outcomes and your Māori customers
- Use and test those mechanisms alongside consultations with your Māori customers

Contributors

AWS contributors to this document include:

- Craig Hind, Solutions Architect, Worldwide Public Sector, ANZ AWS
- Judith Dixon, Public Policy Manager, ANZ AWS
- Paul Keating, Head of Public Policy New Zealand, ANZ AWS
- Ron Amosa, Partner Solutions Architect, WWCO, ANZ AWS
- Bruce Ross, Global Lens Lead for the Well-Architected Framework
- Eddie Gray, Principal Strategic Engagements, ANZ AWS
- Cameron Tod, Senior Solutions Architect, Head of WWPS SA, ANZ AWS
- Adam Barker, Senior Manager, Head of Solutions Architect, AGS, NZ
- Viral Shah, Senior Solutions Architect, AGS, ANZ AWS
- Mat Degerholm, Public Sector Solutions Architect, ANZ AWS
- Tim Woodill, Public Sector Solutions Architect, ANZ AWS
- Roger Somerville, Director, Public Policy, ANZ AWS
- Phoebe St John, Public Policy Manager, ANZ AWS
- Mike Hill, Partner Solutions Architect, WWPS, ANZ AWS

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Māori Data Lens first published.	August 1, 2024

Notices

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.