Mantel group

countdown
woolworths

February 2024

A Guide to Achieving AI at Scale: How Woolworths NZ is Charting their Course

Written by Kathryn Collier, Mantel Group

During this ever-evolving junction in technology, where Artificial Intelligence continually redefines innovation, the imperative to scale AI solutions has never been more pronounced.

Many of our clients are actively exploring Artificial Intelligence (AI), with several already investing in Proof of Value (PoV) initiatives or Minimal Viable Product (MVP) development. However, creating an AI model - whether it's predicting a customer event or a Large Language Model (LLM) - is the easy part. The true challenge lies in establishing an ecosystem for AI across the business that enables long term benefits to be released from AI in a maintainable and sustainable way.

In this whitepaper, we delve into what it takes to achieve AI at scale. We will highlight a recent success story from Woolworths New Zealand, demonstrating how a strategic partnership, along with a focus on people, processes, and technology, has supported a more forward-thinking, data-driven business, better prepared for evolving environments. This case study illustrates how

standardised processes, Machine Learning Operations (MLOps), and responsible AI practices can yield tangible results, empowering organisations to unlock the full potential of AI at scale while ensuring ethical and responsible deployment.

Woolworths NZ recognised they were at a pivotal turning point in their AI maturity when they partnered with Mantel Group. We look forward to sharing this story with you.

Kathryn Collier
Head of Data Science at Mantel Group

"Our wider business is reliant on our Analytics & Insight (A&I) team at Woolworths NZ. In May 2023, we reached an inflection point where servicing the business became blocked by maintaining our existing models or the feature library that supported them. We were bogged down by technical debt in the absence of having a method to maintain models in a light touch, hands off way. Ultimately this hindered the A&I team's capacity to prioritise and perform more strategic, foundational pieces of work, including Personalisation use cases. While so much time was being spent hand holding and maintaining our existing AI capability, we were not able to scale our AI in a way to best support the business."
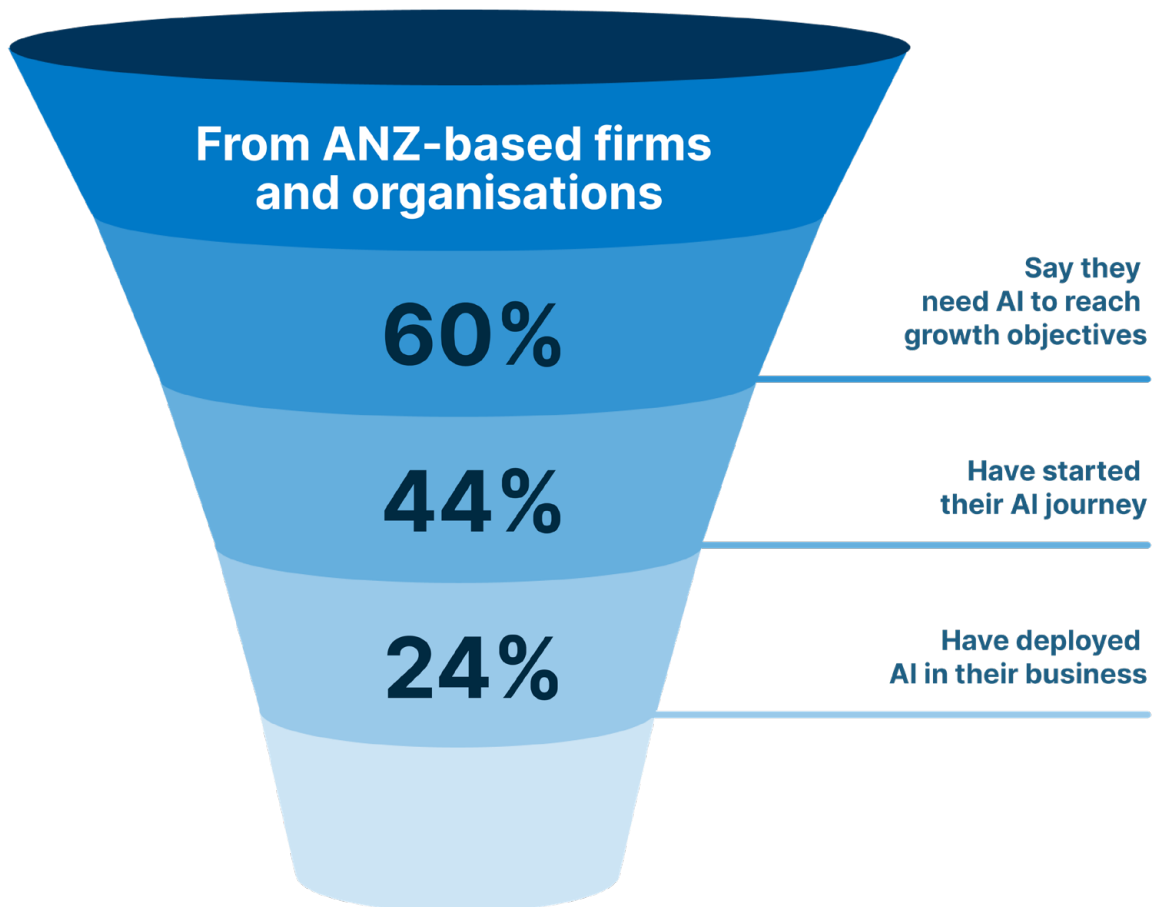
Bonnie Law
Head of Data Science for Woolworths NZ

What is Scalable AI?

In the realm of data-driven organisations, scalability extends beyond the mere capacity to handle increased data and requirements. It fundamentally involves laying the groundwork to ensure that an organisation's people, processes, and technology can sustain the business through phases of growth without necessitating team expansion or additional investments. Scalable AI, in this context, becomes the linchpin of future-proofing data-driven success.

> At its core, scalable AI is about **future-proofing** data science efforts; ensuring that the models and solutions created today can meet the **demands of tomorrow.**

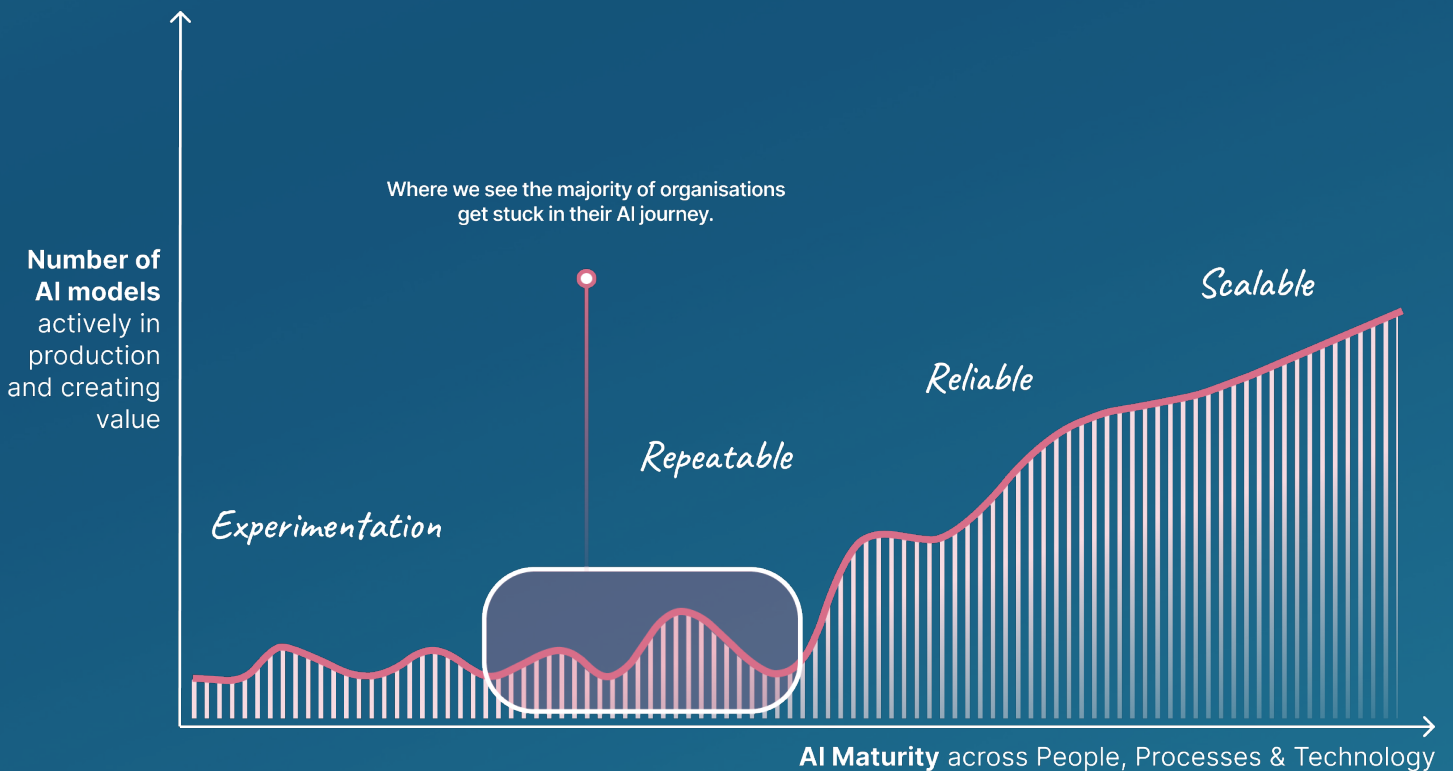**From ANZ-based firms and organisations**

60% Say they need AI to reach growth objectives

44% Have started their AI journey

24% Have deployed AI in their business

Sources: CSIRO AI Ecosytem Report, IBM Global AI Adoption Index 2022

According to a recent report from CSIRO, 60% of Australian-based firms have said that they will need AI to reach their growth objectives. In another report from IBM only 24% of ANZ organisations have been able to deploy AI in their business.

We see a large number of businesses in ANZ struggling to progress with their AI initiatives at scale. There are a number of recurring challenges that exist when organisations are working toward achieving scalable AI solutions. We'll delve into each of these common blockers to scaling with AI and provide a lens on how Woolworths NZ has been navigating these in their journey to achieving an AI ecosystem that enables sustainable growth.

'AI in Business' Maturity Model

Where we see the majority of organisations get stuck in their AI journey.

Number of AI models actively in production and creating value

Experimentation

Repeatable

Reliable

Scalable

AI Maturity across People, Processes & Technology

What does it take to achieve AI at scale?

Developing an impactful AI toolkit, characterised by efficiency, repeatability, reliability and interpretability, without incurring technical debt or necessitating additional resources, represents a formidable undertaking in the pursuit of sustainable and ethical AI solutions.

Establishing an 'AI ecosystem' to ideate, deploy, stress test and maintain AI models sustainably takes a village.

Scaling AI relies on consistent model production to ensure reliable results. Inconsistencies can erode trust in AI-driven insights, leading to incorrect decisions and potential reputational damage. Therefore, organisations must make investments in maintaining precise models.

To address the issue of unreliable AI, the implementation of automation and robust Machine Learning Operations (MLOps) is essential. These practices provide the necessary processes and infrastructure for reproducibility, automated monitoring, and reducing the need for manual interventions and maintenance. This significantly enhances scalability and dependability, thereby simplifying AI deployment with confidence.

Successful AI adoption and growth within organisations can be tethered back to three fundamental themes; people, well-defined processes, and the right-sized technology solutions.

Woolworths NZ discovered that their journey toward impactful development of AI at scale hinged upon elements within each of these three fundamental areas across People, Process and Technology. Woolworths NZ partnered with Mantel Group to support them to carve a clear path forward to sustainable AI practices at scale.

Woolworths NZ recognised the need to improve their approach to ML model development, particularly in terms of standardising processes and practices. Models were often developed in isolation using ML notebooks, leading to an increasing burden of technical debt. As a result, operationalising these models demanded significant productionisation and manual maintenance efforts. Additionally, there was a heavy reliance on individual team members for knowledge.

These challenges are highly common and require unwavering dedication and engagement to address and overcome. In the forthcoming section, we will delve into the process of aligning people, processes, and technology to achieve AI scalability, drawing valuable insights from Woolworth NZ's as they embark on their journey to achieve scalable AI.
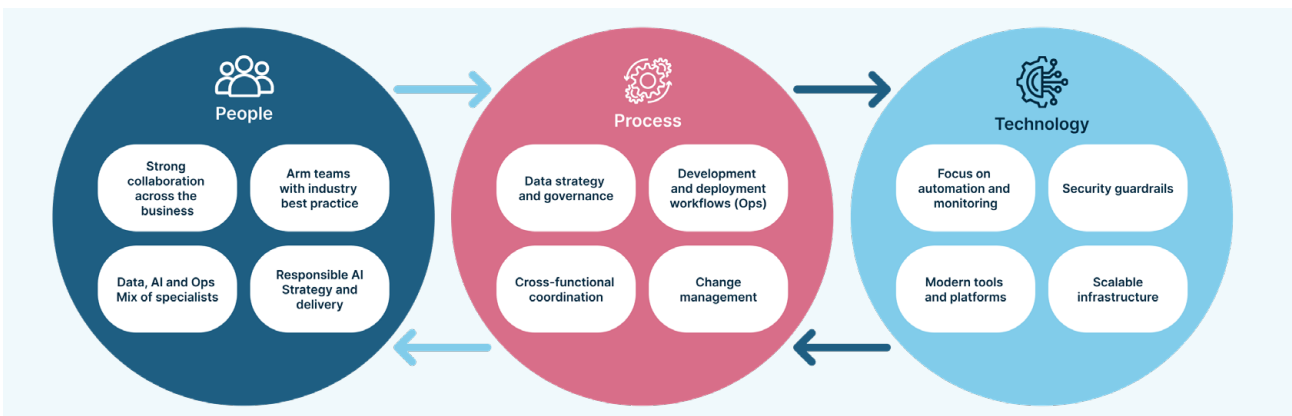
People:
Building the Right Team for Sustainable AI Management

Multi-disciplinary approach to building AI

Scaling AI responsibly takes a village. There is no one team alone that can drive a successful roll out of AI at scale. Successful AI at scale, starts with company wide understanding and appreciation of their role to play in the AI development lifecycle, from experimenting, building, deploying and maintaining an AI model end to end.

Encouraging cross-functional collaboration as well as clearly articulated roles, responsibilities and hand over points will also help enable sustainable scale.

Identifying AI use-cases should be a venture shared by technical and business leaders. A venture which meets somewhere between the top down strategic view of FY planning and the bottom up lens of what is feasible with data available and in a 'fit state'. Engaging business end-users consistently throughout the development process aligns AI solutions with real-world business needs, facilitating seamless integration into existing processes. This inclusive approach ensures that AI solutions not only excel technically but bring value that aligns with the businesses strategic objectives with an output that is being harnessed and utilised effectively.

Diversity within AI teams sparks creativity and **innovation**, leading to robust and ethically sound AI systems.

This approach and open dialogue can also serve to support the business to identify where investment should be focused. For example, if there is a strategic focus on improving customer retention, but the business is not currently capturing the reason customers are leaving - there's an opportunity unearthed that can be addressed.

An open dialogue between the business leaders and technical practitioners of AI is key. It sounds like a simple solution, but it is often the reason that businesses fail to scale with AI. It is also the most common reason that businesses get frustrated at not seeing the ROI in the AI investments, they're too often not consulted or part not of the conversation from the get go.

Dynamic Knowledge Management

Organisations must establish a culture of continuous learning and upskilling to align with the latest AI technologies, tools, and best practices essential for scalable AI solutions. Providing training programs, online courses, and access to AI resources empowers employees to stay current in this dynamic field.

However, offering learning opportunities alone is insufficient. Effective knowledge sharing is paramount, enabling the dissemination of insights, best practices, and lessons learned throughout the organisation. Centralised repositories, comprehensive documentation, and collaborative platforms facilitate seamless knowledge exchange.

Moreover, responsible AI education is integral to this journey. Team members must grasp ethical considerations, potential biases, fairness concerns, and transparency aspects inherent in AI. This knowledge serves as the bedrock for responsible AI practices, ensuring that AI initiatives consistently uphold ethical standards.

Fostering a culture of continuous learning and responsible AI education equips organisations with the knowledge and skills needed to navigate the intricate terrain of AI effectively.

In the case of Woolworths New Zealand, steps had already been established in this space. They had established cross-functional squads in various business streams to support advanced analytics capabilities. While this approach provided agility in delivering value to customers, it posed challenges related to knowledge sharing between squads and in laying the groundwork for standardised practices.

As part of Mantel Group's partnership with Woolworths NZ, we invested in upskilling the team by offering comprehensive training covering end to end aspects of machine learning projects, starting from the initial stages of use case ideation to the long-term maintenance of ML models in production. Collaboratively with the Woolworths NZ team, we put these practices into action by implementing an AI solution for forecasting weekly store sales, concurrently establishing the tools and processes for a novel Machine Learning Foundations framework and a uniform AI workflow.

Additionally, Mantel Group recommended Woolworths NZ to establish a Data Science (DS) and Data Engineering (DE) forum to encourage knowledge and insight exchange among team members across various AI projects. This forum promoted collaborative learning and the sharing of experiences, fostering a culture of continuous improvement, enhancing the efficiency and effectiveness of AI teams, and promoting the standardisation of practices.

Furthermore, the forum provided a platform for exploring cross-squad projects, opening up new opportunities for collaboration and innovation.

Democratisation of knowledge is essential to enable sustainable scale in AI.

The adoption of well established knowledge management also served to release the single points of failure that reside in most, if not all businesses. Scalable AI requires the democratisation of end to end knowledge of the AI model to de-risk that single person or finite group being a dependency to the business's ability to scale, even as the business evolves and the team members change.

Process:
Establishing Governing Principles and Guidance for AI

Scaling AI Responsibly

When we talk about responsible AI, Mantel Group classifies this into three key areas; responsible practices, ethical considerations and awareness of negative externalities of AI.

In the context of scaling AI, all of these elements are of high importance. Responsible practices are arguably the most fundamental area in building scalable AI responsibly for the simple reason that responsibly building AI lends itself to bring interpretability that supports ethical reviews and efficiencies that support a more concise awareness of the unseen costs of AI.

The impact of high computation on the environment is the substantial increase in energy consumption and carbon emissions associated with data centres and supercomputing, contributing to climate change and environmental degradation. Responsible AI remains high on the agenda for Woolworths New Zealand with a strong focus to always do right by their customers and lower their carbon footprint wherever possible. For this reason, in supporting to build a scalable AI capability, we dedicated a large position of our partnership to ensuring responsible practices in how AI was being developed and maintained, end to end.

Woolworths New Zealand will continue to strive towards responsible and repeatable coding practices which lends itself to interpretable code, enabling a more transparent approach to their ethical review.

Responsible practices are a conscious way of working in how you approach building AI. By adopting efficiently coded, repeatable modularised coding practices which are well documented, it will enable you to scale more effectively while also creating an interpretable environment that is not reliant on individuals. It is a time investment not only to set up and establish processes around it, but also in the upskilling and application of the responsible coding practices but it is an investment that will pay dividends when there is a raft of AI models in production that are all easily interpretable and operating effectively with little to no intervention.

Throughout the engagement, Mantel Group established knowledge management coupled with training on efficient and effective coding practices to support Woolworths NZ to achieve a uniform approach to building models while consciously minimising unnecessary compute at scale where possible.
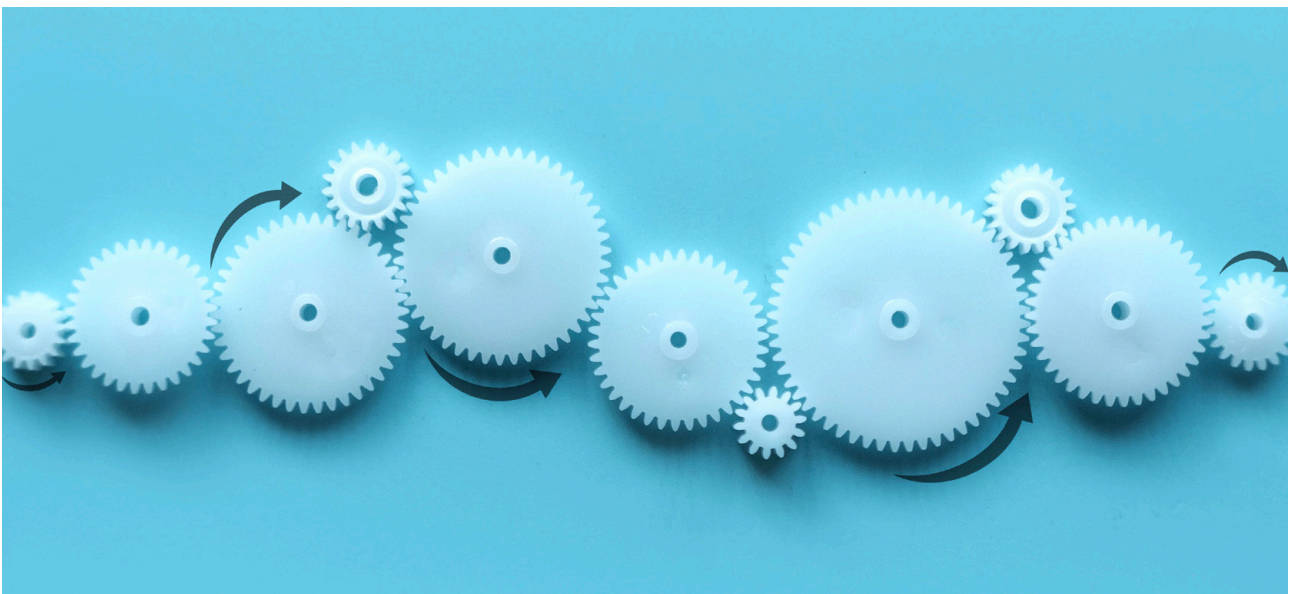
Establishing a Uniform AI Workflow

Standardising processes to ensure consistency and promote best (responsible) practices in model development and deployment is what supports AI to scale sustainably for the long term. It is not enough to simply change the way in which models are developed within the technical team. In order to establish long term adoption, a clear understanding of the how and why it matters is required across the business.

Without end to end business buy-in, best efforts to sustain scalable and responsible AI practices will be weathered by the pressure for results and heady timelines.

Establishing clear ownership, roles and responsibilities across the business (e.g. which role in the business is responsible for monitoring this model ongoing) will support to drive this adoption.

Mantel Group supported Woolworths NZ to establish both the technical workflows as well as a broader understanding of the model development process for the business lens to support long term adoption of scalable AI. This took the form of brown bags covering topics such as; *How to Manage a Model Long Term*.

In addition to the establishment of Machine Learning Foundations, a centralised set of comprehensive approaches and guides to building ML models end-to-end was developed, while incorporating best practice methodologies. This documentation served as a valuable resource for both existing and new team members, ensuring that they had a clear roadmap for developing AI solutions. By sharing these best practices, Mantel Group developed the 'Woolworths NZ Playbook' for maintaining consistently high quality and uniformity in the execution of AI projects.

Mantel Group collaborated with Woolworths NZ's A&I Delivery team to design and roll out in depth training for those responsible for AI project delivery. Success in AI projects hinges on their leadership, stakeholder management and the ability to capture requirements effectively. This training equips the delivery team at Woolworths NZ with the essential skills to work closely with stakeholders, define project requirements and balance technical feasibility with real-world applicability. Through this training investment, Woolworths NZ is supporting a team of project leaders to drive AI projects towards success, aligning them with their strategic vision and playing an essential translation role between the business outcome and the technical application.

Woolworths NZ is demonstrating a commitment to excellence with the initiation of the MLOps (Machine Learning Operations) squad. This specialised team consists of individuals with an interest in MLOps. This team aims to provide support and enhancements for the newly established Machine Learning Foundation framework. The ambition of this squad is to champion best practice approaches of their AI solutions, to bring them closer to their long-term goal of sustainable AI at scale as they continue to develop and mature their AI capability.

Planting the seed for change

Integrating AI requires a change management strategy to understand stakeholder expectations and navigate any resistance to change. Inherent mistrust in AI, often stemming from concerns about its decision-making processes, bias, and lack of transparency, can significantly limit the scale and adoption of AI.

Risk of distrust of AI solutions can creep in where lack of understanding and transparency exists.

This distrust can hinder organisations from fully realising the potential benefits of AI, such as improved efficiency, data-driven insights and automation of repetitive tasks. To overcome this challenge, it is crucial to build transparency, accountability, and fairness into AI systems, and to educate stakeholders about how AI works and its potential to enhance human capabilities rather than replace them. Building trust in AI is not only an ethical imperative but also a key factor for its successful integration and scalability.

The squad based approach that Woolworths NZ adopts lends itself well to a better connection between the business and the development of AI solutions, where change and agility is part of their way of working. As part of the delivery training, a key focus was on understanding all of the touch points and impact areas of the AI capability, such that timely and effective communication could seed an understanding of what lay ahead once the solution was ready. Not only did this allow the stakeholders to play a role to input into the solution, it also set the expectation that change was afoot and they were part of making that change as impactful as possible as the SME (Subject Matter Expert). In general, having this collaboration and input up front often plays a pivotal role for

achieving the buy-in and wider adoption of the AI solutions developed, not to mention leading to more impactful outcomes. SME input on integration requirements, security thresholds, resource allocation requirements etc. will also open up planning to ensure a successful outcome at scale with a broader understanding beyond just the statistical methodology required to build it.

The caveat to this approach however is the invitation to a cast of thousands to give input, which can have a detrimental impact on progress (not to mention stress levels). Initial group workshops can offer the opportunity for voices to be heard in a single space and allow understanding across all stakeholders at once to be realised. Another effective tool in your arsenal when it comes to stakeholder management is stakeholder mapping.

Stakeholder Map: Who needs what?
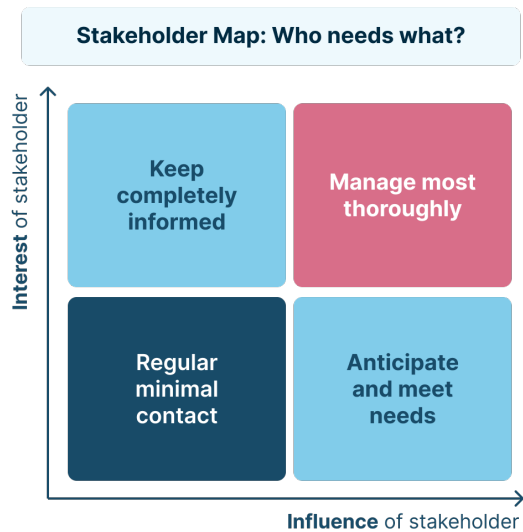
This allows you to weigh the significance and importance of the input shared and enable a more streamlined process to absorb information without losing project momentum.

Data and Operational Governance

Effective governance is imperative for ensuring data and model quality at scale, as well as security and compliance. The absence of robust governance can lead to inconsistencies, privacy breaches, and regulatory risks, impacting decision-making. Siloed data further hampers scalability and complicates data and AI integration.

Technical debt is an **anchor** that hampers scalability and innovation.

Breaking Free of Technical Debt to Enable Scale

Technical debt arises from factors like deadline pressure, resource constraints, changing requirements, inadequate documentation, limited experience, legacy systems, unforeseen issues, neglected maintenance, scope creep, and lack of code reviews, necessitating a proactive approach to balance short-term goals with long-term sustainability.

Accumulating technical debt has a long-term cost of maintenance, not to mention an impact on team morale. Outdated data infrastructure only aggravates the issue. Addressing technical debt requires a deliberate strategy. While it may be the technical practitioners that are tasked with "fixing" the issues, it is by no means solely their responsibility to ensure that it doesn't happen again. This takes a company wide approach to ensure that technical practitioners are not in a position where quick fixes are the only way to meet the businesses expectations.

By introducing solid knowledge management, a better understanding and engagement from the business and responsible coding practices, organisations can start to chip away at repairing technical debt and ensure that guardrails are in place to prevent it occuring again. This proactive approach unburdens organisations, paving the way for agile, efficient, and scalable AI solutions.

Technology: The Foundation of Scalable AI

Along with people and process, technology forms the third pillar crucial for establishing the foundations of Scalable AI by providing the tools, infrastructure and methodologies needed to develop, deploy and manage AI systems at scale.

Scalable Infrastructure and Cloud-Based Platforms

Scalable infrastructure, advanced AI development frameworks, and cloud-based platforms collectively empower organisations to effectively handle increasing workloads and manage the growing demand and data volume.

Cloud-based infrastructure plays an important role in achieving scalability, providing the flexibility needed for resource allocation. Leveraging cutting-edge AI development frameworks and cloud-based platforms, complete with pre-built components, accelerates the development process, ensuring that AI solutions are not only reliable but also easily repeatable and monitored during deployment.

To further enhance the efficiency of AI deployment, organisations are increasingly adopting DevOps and MLOps practices. These practices emphasise automation, collaboration, and continuous integration and delivery, reducing deployment time and increasing the overall reliability of AI solutions. Containerization and orchestration technologies, such as Docker and Kubernetes further simplify the deployment and create consistent environments for AI applications.

This foundation accommodates growth, accelerates development, streamlines deployment, and enhances overall solution reliability.

Mantel Group supported Woolworths NZ in implementing these foundations by utilising Google Cloud's Vertex AI Platform. Vertex AI Platform provides a versatile set of tools that aim to support data scientists and ML engineers throughout their journey in AI development. These tools facilitate the entire Machine Learning workflow, from training and tuning models to deployment and monitoring models long term.

Integrating AI into Business Operations

As organisations increasingly recognise the transformative potential of AI, the seamless integration of artificial intelligence into existing business operations becomes a strategic imperative. This integration involves a meticulous and well-thought-out strategy, as AI solutions must not only coexist but blend with other systems, share data effectively, and, most importantly, deliver tangible value within real-world business contexts.

The key focus areas when integrating AI into business operations include:

- **Interoperability:** Ensure AI systems seamlessly work with existing tech stacks.

- **Data Synergy:** Ensure access to compatible, high-quality data.

- **User Experience:** Prioritise user-friendliness and workflow continuity.

- **Real-World Value:** Align AI with core business goals and measure its impact through KPIs.

- **Scalability and Adaptability:** Ensure that AI solutions can grow and evolve with changing needs.

- **Continuous Improvement:** Establish feedback loops and iterative refinement.

- **Change Management and Training:** Train employees and manage change effectively for smooth adoption.

By addressing these key points, organisations can harness AI's transformative potential and drive innovation within their operations.

Rapid Prototyping and Standardisation

Rapid prototyping involves the swift creation and testing of AI models to accelerate development and ensure adaptability for real-world scenarios. It is essential for enabling quick experimentation and development acceleration. Efficiently transitioning from prototypes to production is equally critical. This transition is vital for demonstrating ROI, seamless integration, and automation.

For Woolworths NZ, we established a fully functioning end-to-end Machine Learning capability, incorporating Mantel Group's best-in-class MLOps practices. This included creating reliable and repeatable code and pipeline artefacts that will live on in Woolworth NZ AI journey as they continue to scale their capability.

Key Components of the Solution

- **ML Repository Template**
We introduced a structured ML repository template with reusable code to simplify new model prototyping, demonstrate best practice DevOps, and facilitate peer code reviews. This template repository forms the starting point for any new AI use case development ensuring a standardised and accelerated development.

- **Model Experiment Tracking**
We set up experiment tracking, a practice of systematically recording and monitoring experiments conducted during the development and training of machine learning models. This plays a significant role in ensuring model reproducibility, transparency, and effective collaboration within data science teams.

- **Data Analytics and Model Development Strategies**
Our approach covered data quality assurance, exploratory data analysis, hyperparameter tuning, and performance evaluation.

- **Automated and Repeatable Pipelines**
We implemented automated pipelines for model training, evaluation, inference, and monitoring. These pipelines covered automated unit tests helping catch issues early in the development process and ensure good code quality.

- **Model Registry**
A model registry was introduced to track model versions and evaluations, simplifying model deployment for batch predictions and online endpoints.

- **CI/CD Workflow**
We integrated CI/CD workflows with automated repeatable steps to establish a well-defined code promotion pipeline that includes multiple stages such as development, testing, staging, and production. This ensures that code changes are thoroughly tested before being deployed to production.

- **Monitoring and Alerting**
We implemented monitoring and alerting systems for deployed models, enabling easy tracking of performance, prediction drift, and training serving skew.

> **Woolworths NZ are now equipped with the foundational tools to establish an efficient machine learning capability with process guides, trained team members and reusable templates. They are stepping away from a reactive model to a more proactive, scalable AI approach, where the organisation is better-prepared for future data-driven challenges.**
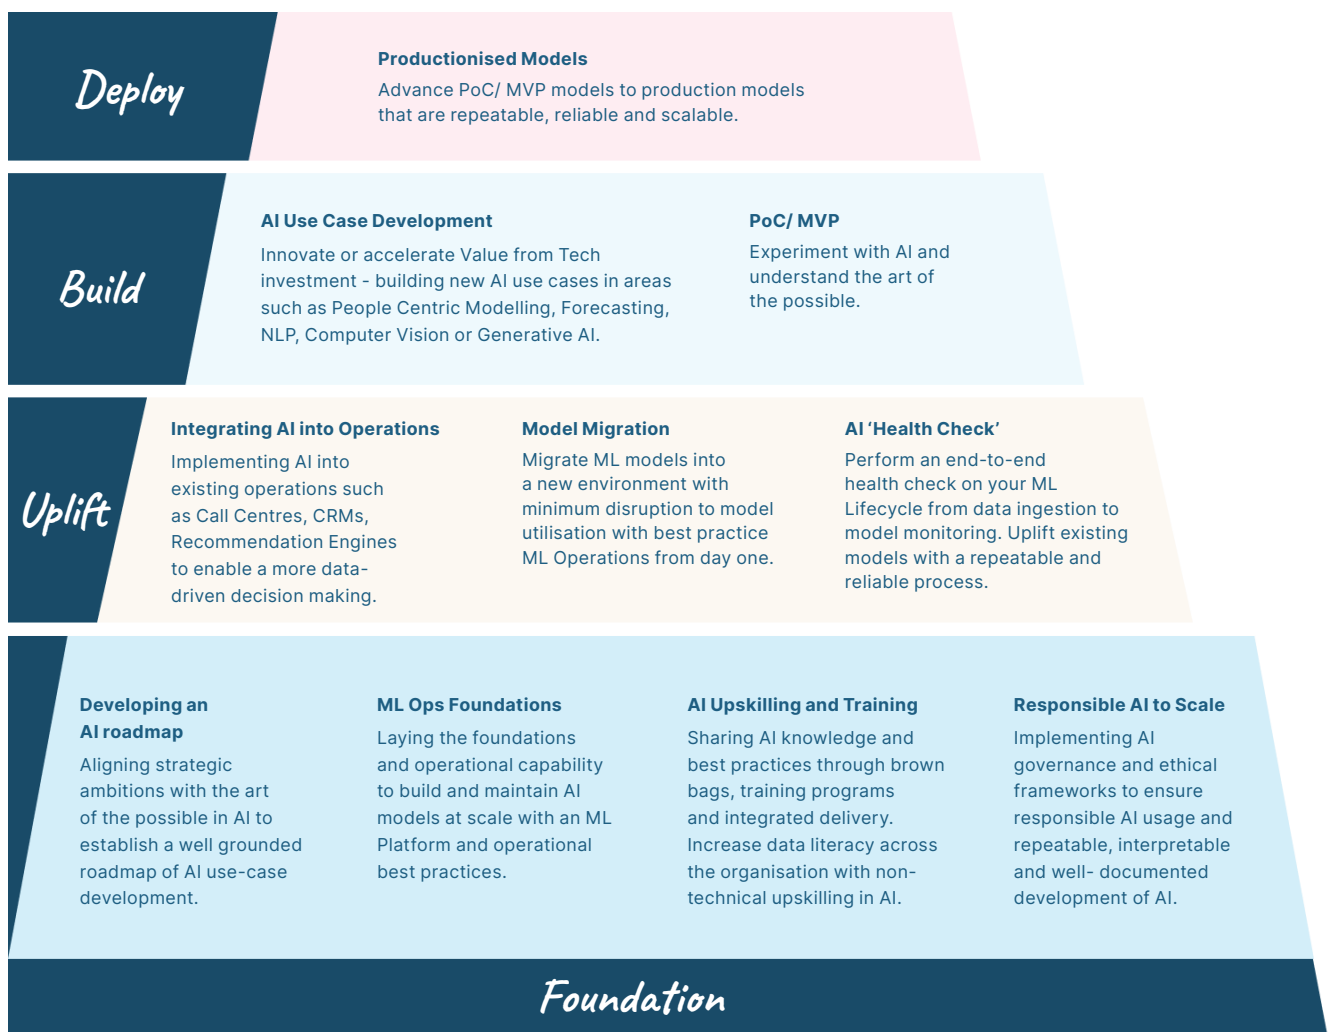
What Scalable AI Enables Businesses to Achieve

Woolworths NZ achieved remarkable progress in just eight months with our small team of three to four full-time employees. They've entered a new era of ML capabilities and a better understanding of AI across both technical and business teams. While still at the beginning of their journey, Woolworths NZ fosters a determination to strive towards wider adoption of responsible AI practices that drive tangible results at scale.

Next Steps for Scaling AI in Your Business

In order to embark on a journey of scaling AI, organisations need a clear roadmap that ensures success is achievable. The initial step involves evaluating your organisation's readiness for AI by determining the current AI maturity level. This entails assessing current AI capabilities while simultaneously identifying areas for improvement.

To make scalable AI a reality, organisations should integrate AI delivery seamlessly into their existing processes, ensuring alignment with broader business objectives. This holistic approach enables organisations to leverage AI effectively, fostering innovation and growth while keeping pace with the dynamic AI landscape. With the right strategy and approach, scalable AI is achievable for organisations of all sizes and sectors.

The below figure summarises the key themes of this whitepaper and demonstrates how an organisation can work toward Scalable AI.

Deploy

Productionised Models
Advance PoC/ MVP models to production models that are repeatable, reliable and scalable.

Build

AI Use Case Development
Innovate or accelerate Value from Tech investment - building new AI use cases in areas such as People Centric Modelling, Forecasting, NLP, Computer Vision or Generative AI.

PoC/ MVP
Experiment with AI and understand the art of the possible.

Uplift

Integrating AI into Operations
Implementing AI into existing operations such as Call Centres, CRMs, Recommendation Engines to enable a more data-driven decision making.

Model Migration
Migrate ML models into a new environment with minimum disruption to model utilisation with best practice ML Operations from day one.

AI 'Health Check'
Perform an end-to-end health check on your ML Lifecycle from data ingestion to model monitoring. Uplift existing models with a repeatable and reliable process.

Developing an AI roadmap
Aligning strategic ambitions with the art of the possible in AI to establish a well grounded roadmap of AI use-case development.

ML Ops Foundations
Laying the foundations and operational capability to build and maintain AI models at scale with an ML Platform and operational best practices.

AI Upskilling and Training
Sharing AI knowledge and best practices through brown bags, training programs and integrated delivery. Increase data literacy across the organisation with non-technical upskilling in AI.

Responsible AI to Scale
Implementing AI governance and ethical frameworks to ensure responsible AI usage and repeatable, interpretable and well- documented development of AI.

Foundation

Final Thoughts

Scaling AI requires a holistic approach which incorporates responsible AI practices, strong business foundations, unified technical teams and a well-appointed ML platform.

Scaling AI requires a holistic approach which incorporates responsible AI practices, strong business foundations, unified technical teams and a well-appointed ML platform. By navigating these challenges, businesses can innovate while maintaining ethical and transparent AI practices. This journey toward scalable AI not only ensures regulatory compliance but also establishes a trusted and innovative presence.

In the age of data-driven decision-making, organisations that tackle the challenges which hinder scalable AI with a holistic approach can unlock transformative capabilities. By fostering cross-disciplinary teams, implementing robust governance and adopting the right technology infrastructure, businesses can thrive in an increasingly AI-driven world, achieving agility, innovation and competitiveness.

At Mantel Group we're always keen to start new conversations on using technology to impact people in a positive way.

Reach out today

Speak with Kathryn Collier about how your business can achieve an AI capability that scales.

✉ kathryn.collier@mantelgroup.com.au

Mantel group

Unlock your AI potential

We're always keen to start new conversations on using technology to impact people in a positive way.

Contact Us ↗