How to Build a Rock-Solid AI Strategy

mesh-ai.com

Table of Contents

Michael Chalmers

CEO

Michael is a customer centric, outcome focused business leader with a successful track record in driving growth, who is proud of building and leading diverse, high performing, cross functional global teams.

As the CEO of Mesh-AI, Michael is responsible for the continued growth of the business across all functions, leading end-to-end business operations. Since joining Mesh-AI in January 2023, Michael's focus has been on delivering measurable, market leading outcomes for our highly regulated enterprise customers, specialising in data and AI transformation.

Michael has over 15 years of experience working with enterprise customers to transform them into digitally enabled organisations. Prior to joining Mesh-AI, Michael led Contino, a global devops and cloud consultancy as Managing Director, before that he held senior positions at Hewlett Packard Enterprise as a Chief Technologist and Global Account Director, as well as stints with Fujitsu, SecureData and HP - steadfast as a technical evangelist for enterprise clients.

Chapter 1

Introduction

AI investments aren't delivering.

Despite pouring massive resources into AI initiatives, they aren't transforming businesses.

Why not?

The problem is that businesses are not treating AI with enough *strategic* depth.

Typically, AI is developed in isolation to wider business operations and strategy, often treated as "garage" type development and built around POC's. This means, even if these projects are successful, converting them into long lasting enterprise wide deployments is rarely achieved.

POCs and experiments are not enough to transform the enterprise.

It's not that AI doesn't have incredible potential. The way ChatGPT has recently captured the public imagination and the impressive outputs it can deliver attest to the potency of the technology.

But that capability needs to be carefully nurtured, given the nutrients it needs to grow and your business structure and culture shaped to enable it to flourish.

This is analogous to other game-changing technologies. Take the Internet, you can't just 'do Internet' and create business value. You still need to be strategic, creative and bold with how you use it, and prepare for any pending bubbles that leave those who haven't in the dust.

And for AI to truly deliver on its potential, it must be deeply rooted at the core of your business, supported by sturdy technical foundations and held within a business culture that appreciates its value and is educated to be able to use it.

In short: you need to deliver a deep-rooted, rock-solid AI strategy that puts AI at the heart of your business.

In this eBook, we'll cover the **5 pillars** that you need to build an AI strategy that delivers real ROI in a sustainable and scalable way.

Chapter 2

The Business Case For AI

Fuelled by the rise of deep learning, compute accessibility and the open source community, AI has become a real opportunity for achieving true business impact.

Take Amazon, for example. The idea of an online e-commerce store wasn't revolutionary. But how they used AI to innovate and optimise their entire infrastructure and delivery model is. These foundations also set them up as a business to consistently innovate, improve and ultimately win.

AI is ultimately an enabler. It can be used to automate simple repetitive tasks but it can also surface connections and optimise processes, in ways humans could never have imagined. You must build your business around both AI and humans, understanding where to deploy together and ultimately enable us to deliver more value.

Here are 5 key ways that AI can add value to businesses:

1

Improve productivity and operational efficiency

This can include evaluating resource utilisation, analysing historic inventory management, and delivery and distribution

2

Automate tasks

Relieve humans of lower-value repetitive tasks freeing them to do more creative and valuable work

3

Developer new product and service ideas

Leverage your data to identify trends, patterns and ideas that can be turned into opportunities to develop new products and services

4

Improve and augment decision making

AI systems are able to support decisions through real-time and up-to-date data gathering, forecasting, and trend analysis

5

Personalise the customer experience

Analyse data to understand your customer, who they are, their needs to deliver personalised journeys and product offerings

But it's very easy to do AI poorly and miss out on these benefits.

Chapter 3

The 5 Pillars of AI

We have broken down a rock-solid AI strategy into 5 core pillars that covers the entire AI value chain across people, process and technology.

01 Culture

02 Ideation

03 Delivery

04 Trust

05 Impact

Let's explore each of these in greater detail.

1. Culture

A successful AI culture provides the nourishing long-term strategic context and environment in which your AI initiatives can grow and mature

What It Is

An AI culture is one that deeply recognises the business value of AI and makes it a core part of company vision, strategy and communication.

Why It Matters

An effective culture puts AI at the heart of your business, with people at all levels of the business well-informed about its potential applications and business value .

Without this strong cultural foundation, the long-term, end-to-end context is lost and AI is reduced to throwing short-term technical solutions at the wall to see what sticks.

Done well, a good AI Culture generates trust in the technology, helps everyone understand the benefits of AI initiatives and ensures that it is executed in a well-thought-out, holistic, end-to-end way that is clearly aligned with business goals

Key Principles

Leadership

Your leadership team must deeply grasp the potential impact AI could have within your business and proactively promote AI as part of the wider business strategy. They must consolidate and disseminate a consistent vision for AI success to the rest of the business, while ensuring that they have the right people in the right place with the right technical implementation plan.

Strategy

You need a long-term end-to-end strategy in place before you make significant investments.

You can't focus only on 'innovation' or deploying models but must consider the entire value chain, from the first ideas to the impact on the end user or business.

The strategy helps everything you do to be aligned with and guided by overarching business goals and operating model.

The ideal AI strategy is one that should be continuously reviewed based on technology advancements, validated learning and shifting business priorities.

Education and Literacy

Education is critical to help people grasp the art of the possible with AI, how it can help them and your business for the better. Without this, they may be mistrustful of the technology or feel that their job could be threatened. It's critical to help them understand AI as a powerful business tool and a valuable support for their work.

With an AI-literate workforce, you can take your whole business on the AI journey with you. They can then help to spot opportunities for AI within the business.

Key Questions

1. Do you know what your business goals are and is there a clear strategy for how the business can use AI to achieve them?

2. Is your leadership team communicating a consistent vision for AI as a core part of the wider business strategy?

3. Is the strategy widely understood and synergistically aligned with other business functions?

4. Do you have the right mechanisms to democratise knowledge on AI and your strategy?

5. Do you have the right experts within the business and are you connecting them with other teams and functions?

2. Ideation

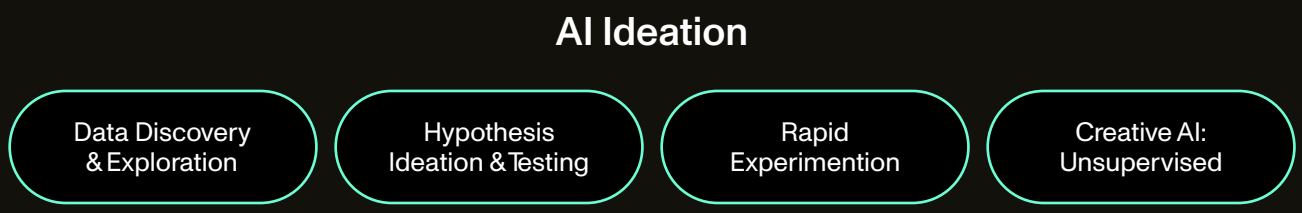

Develop creative, relevant and valuable ideas in an iterative and incremental way at the intersection of data science and business expertise.

What It Is

Ideation is the process of identifying use cases for AI that are both highly valuable and technologically feasible.

Teams of data scientists and business experts collaborate to develop and then test hypotheses through a process of rapid experimentation. The results are fed back to the business and the hypotheses are further iterated and refined. This process we call continuous validated learning.

The result is a prioritised list of business-valuable use cases to productionise and scale.

Why It Matters

Using data and AI is an extremely creative process. Amidst nearly infinite possible use cases, developing well designed and valuable ideas and approaches for using AI generates astronomically better returns than just throwing a large number of ML models at a problem.

If the business isn't ideating effectively, you will fail to recgonise valuable opportunities or those AI initiatives you do develop will be sub-optimal. The more you can connect cross organisational ideas, the wider the impact and higher likelihood you will genuinely create a competitive advantage.

Key Principles

Data Discovery and Exploration

Your data scientists need to be able to freely discover and explore your business (and third party) data to develop ideas on how they can extract value from it using AI.

Your data scientists' time is precious. You don't want them spending their time managing infrastructure, struggling to find out what data is available and how they can access it and so on.

Your technology and governance should be such that they can focus as much as possible on what data they want to use, developing and testing hypotheses.

Hypothesis and Ideation Testing

Ideas must be cross-organisational. A co-creative collaboration between the scientists (who know what is possible) and the business (who know what is needed) is ideal.

Similarly, it's critical to have a broad knowledge of what the different types of model can do with different types of data. The more diverse your team is in terms of their AI experience, the more creative they can be with their approaches.

Rapid Experimentation

Once hypotheses have been developed your teams and technology need to be arranged such that you can rapidly test and iterate on them. Data Science and Machine learning workbench environments, with seamless data connections, compute accessibility and collaboration tools are essential to allowing rapid experimentation.

Critical here is a systematic and democratic process for quantifying the impact of potential use cases so that they can be usefully prioritised according to business value. Impact can come in many forms, but the process of outlining performance indicators and "measures of success" are key. The faster you can experiment, the more quickly you will develop valuable use cases.

Unsupervised AI

Supervised AI is when you know what you want to learn and are inputting labelled data into your algorithm. Unsupervised AI is the opposite, when you have lots of data but you are not supervising the outcome. Instead, you leave it up to your algorithms to extract patterns and derive insights in an open way. This can result in very creative and unexpected ideas, hypotheses and use cases to evolve.

Key Questions

1. Are your use cases guided by their business value, not technical qualities?

2. Is there a democratic process for identifying AI use cases and measuring the required effort versus their potential impact?

3. Are business and data experts collaborating across the organisation to identify ideas to put into the experimentation backlog?

4. Is there a 'fail fast' attitude towards developing use cases?

5. Is there a technical foundation for collaboration? Including, for example, cloud collaboration tools, a central repository of reuseable assets/features, up-to-date research on algorithm methodologies etc.

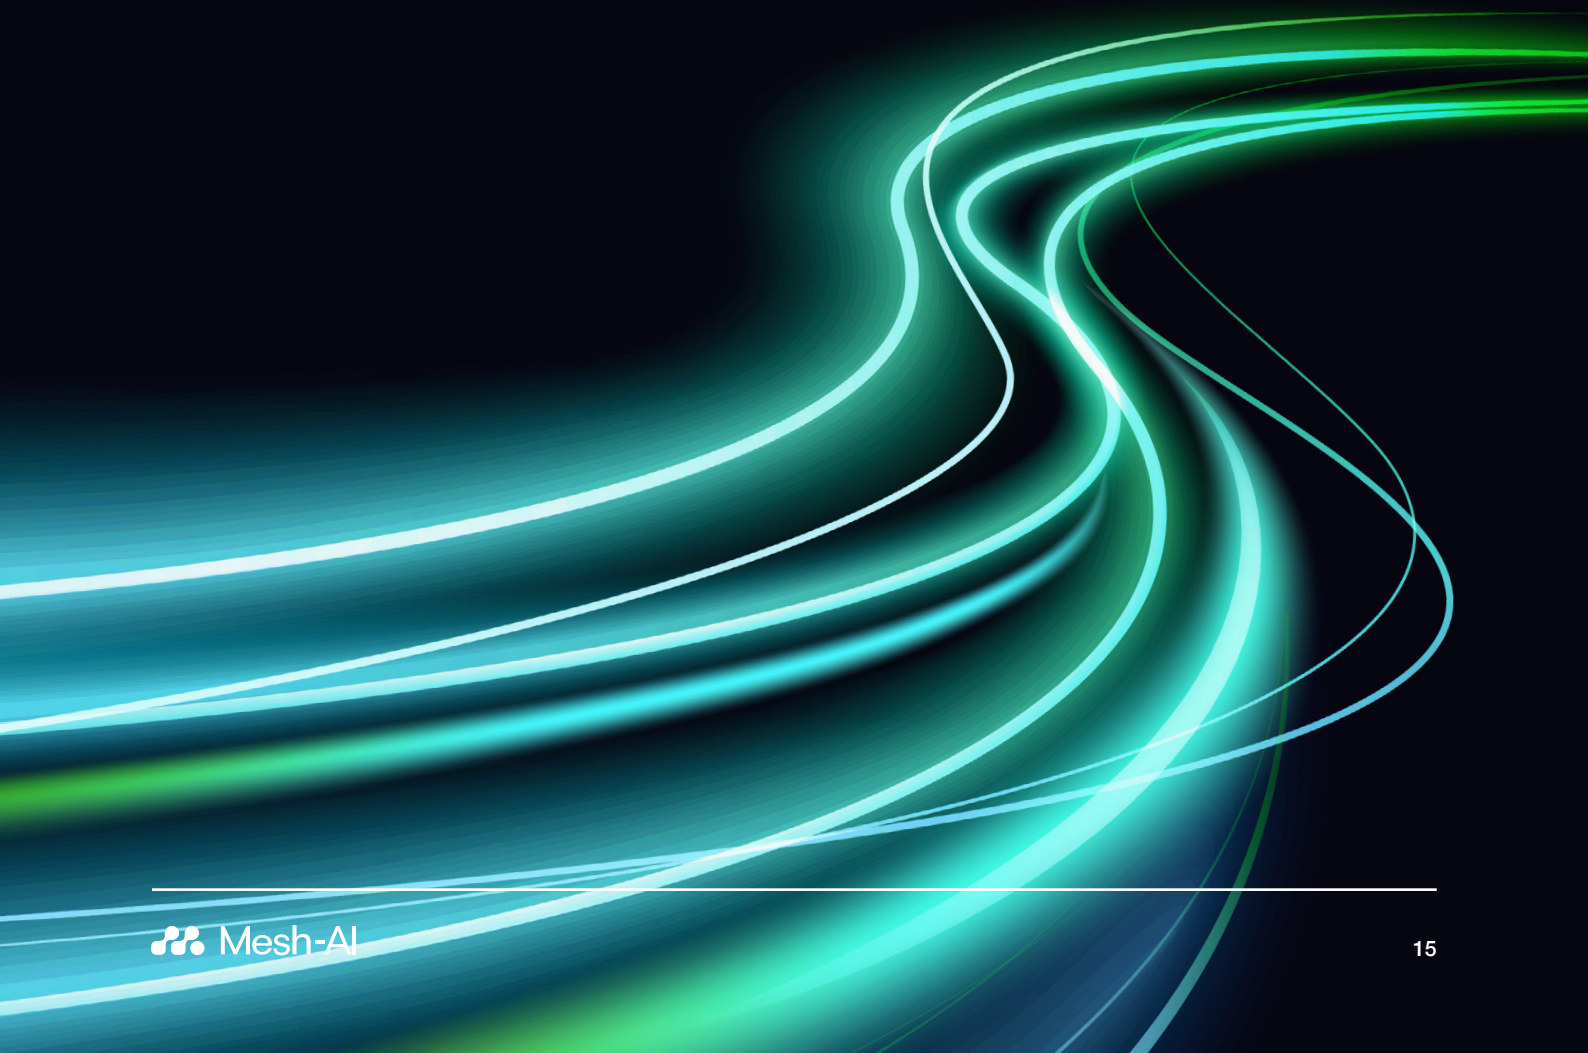

3. Delivery

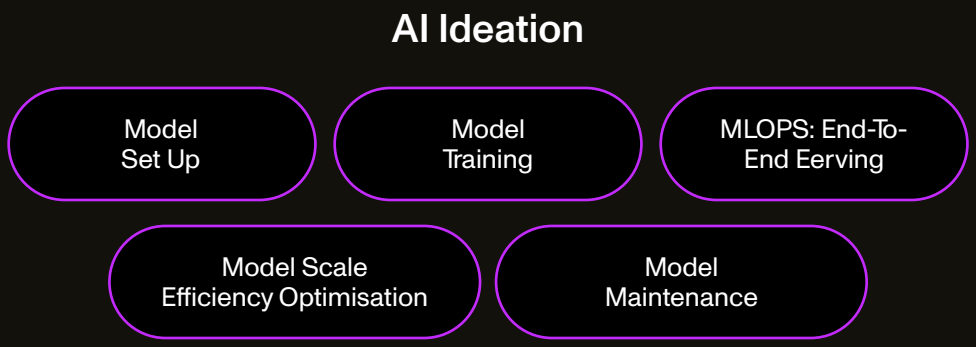

The technology, processes and people to train, deliver and maintain your AI models at scale

What It Is

AI Delivery encompasses the end-to-end process of setting up, training, measuring, scaling up and maintaining your models over time in service of business goals.

This includes technology that can reliably and securely productionise models at scale, processes that govern the end-to-end value chain and teams whose structure and expertise mirrors the delivery lifecycle.

Why It Matters

Without a proper delivery strategy and framework, your AI initiatives will never reach the volume and scale required to be truly impactful. They will never go beyond the experimental stage and they will not evolve with your business or the world as it changes.

AI is like a product, you can't just launch it and then rest on your laurels. Even a highly-accurate model will become irrelevant in a few months. You need to make continuous AI innovation and adjustment part of the fabric of your business.

Even with the plethora of "self serve ML workbenches" available, having a strong strategy in place is key to both selecting and optimising the use of these off the shelf solutions.

Key Principles

Model Setup

You need strong AI development standards and policies in place to support rapid and repeatable model development.

Supportive technology functions like data preprocessing and centralised feature engineering or feature stores mean devs can spend less time on preparatory tasks and more on model development and exploration.

There should also be a built in practice and tech stack that facilitates continuous contribution to shared repositories, approaches and learnings. Ensuring each exploration has the ability to inform the next.

Model Data

Data quality, consistency and reliability is the key when it comes to model training. You need a standardised framework for data management and versioning.

Mitigate ethical risks by having a well-defined training data assessment framework to ensure that balance is maintained and bias avoided.

Model Training

The sheer number of models and frameworks available within machine learning nowadays is massive. Having the knowledge, ability and accessibility of training these models within your team and stack is one thing, but enabling experimentation at scale is another. Your development environment should facilitate rapid model training and experiment tracking.

Like most of the AI space, this is a moving goal post and ensuring your team are tuned in to both algorithmic and technological advancements is key to the models and outputs remaining “state of the art”

MLOps

The term MLOps has become widely adopted as a term but is typically poorly defined as a practice. This should cover the end to end development, deployment and monitoring process, but to make it more actionable we break this up into MLOps for Organisational Management and MLOps for engineering development.

MLOps for Organisational Management

You need a continuous, end-to-end operating model that governs the delivery of AI models, ensuring that they can be seamlessly deployed, carefully optimised and well maintained.

This includes clearly-defined roles and responsibilities for the entire delivery lifecycle, dedicated infrastructure, continuous delivery and a clear governance framework for model delivery.

MLOps needs to be incorporated with the wider operating model and development practices of your wider business, this includes things such as privacy and product management.

MLOps for Engineering Development

Model management, continuous integration continuous deployment (CICD), logging, registry, integration and evaluation frameworks are all required components of a robust MLOps framework and process. Not forgetting the model training itself, which requires scalable compute, algorithm framework accessibility and experimentation procedures.

Model Optimisation

Set clear KPIs to measure the effectiveness and efficiency of models in production and establish a feedback loop mechanism that enables us to incorporate user feedback into the models and optimise them accordingly.

Model Maintenance

You need processes in place to ensure that model performance and security is upheld over time. Your platform should provide automated monitoring and alerting capabilities to detect issues such as data drift, model degradation, and resource usage spikes.

Key Questions

1. Do you have strong AI development standards and policies in place?

2. Are your platforms and pipelines sufficiently automated to allow your data scientists to focus on machine learning rather than infrastructure?

3. Are your data science and machine learning teams able to run and track experiments at scale across a state of the art models and frameworks?

4. Is there a clear data governance framework that covers things like data security, versioning, lineage so that data is consistent, reliable and trustworthy?

5. Are your models in line with clear business goals that are being tracked using appropriate KPIs?

6. Are there clear roles and responsibilities for the Entire AI Value chain?

4. Trust

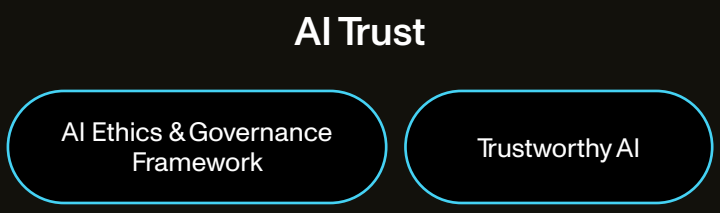

Building explainable, transparent and ethical AI solutions that respect laws, regulations and are grounded in human rights.

What It Is

A trustworthy AI solution is one that complies with relevant laws and regulations, carefully considers ethical risks and biases and the workings of which are transparent.

Governance must be front and centre to ensure trust must be embedded in all AI processes, from initial risk assessments to designing and developing solutions.

Why It Matters

One of the biggest potential barriers to AI adoption is lack of trust. Your people, your customers and regulators need to feel that your technology is safe, transparent and accountable in order for it to be widely adopted.

Trust must be required to ensure that you maximise the business value of your AI initiatives while minimising ethical and legal risks.

Key Principles

Ethics and Governance Framework

As AI becomes more ubiquitous, there is already a growing number of regulatory policies around AI governance and ethics. Meeting these standards by default is an unavoidable and indispensable pillar of any successful AI initiative.

It's important to create a well-defined stance on AI ethics at the organisational level that can be translated into technical requirements. Given the rapid pace of developments in this area it also needs to evolve over time and so embedding as a continuous process is key.

Finally from a procedural perspective, it's important to ensure you have the right guidelines and frameworks in place to detect and enforce these ethical and governance practices.

AI Explainability

AI algorithms are often perceived as "black boxes" whose inner workings are obscure.

'Explainability' refers to an AI system that can answer why it made particular predictions and not others, which helps users to understand it as well as to govern it for appropriate use regarding, for example, gender and race. This is also a strict consideration for certain industries where there are legal rights and protections surround the deployment and usage of AI.

Depending on the level of requirements, your business can choose models with high levels of explainability, or apply downstream algorithms to extract explainability that are more likely to gain trust.

The concept of fairness largely centres around the usage of how machine learning is deployed, however, this is deeply rooted back to the algorithmic design and the data on which it was trained. Having a robust and well enforced set of guidelines and development practices are key to building and tracking the fairness of your AI solutions.

Risk and vulnerabilities of AI are constantly evolving and so building a culture of continuos feedback and improvement is imperative, putting user feedback at the heart of the iterative development process.

Key Questions

1. Are the teams developing AI solutions well-versed in Explainable AI (XAI) techniques?

2. Is your level of risk tolerance clear and well-defined?

3. Do you have a framework for the detection and mitigation of different bias categories (e.g. sample, prejudice, measurement) across the entire AI lifecycle?

4. Is there a mechanism to allow stakeholders to provide feedback on AI governance?

5. Is the developmental lineage of your AI solution (i.e. the data it was trained on, how it was trained) fully traceable and available?

5. Impact

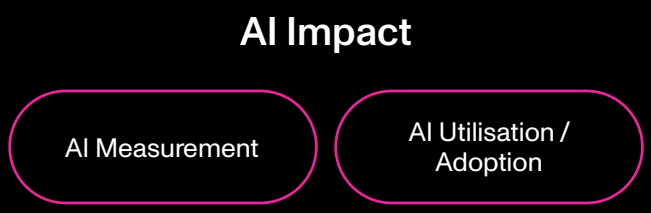

Don't do AI for AI's sake - the impact on the business is what matters most

What It Is

Impact involves measuring and tracking your AI models across a range of key metrics, including performance, adoption and ROI. The goal is to determine how they have impacted business outcomes.

Why It Matters

Many businesses simply do not know what impact their AI initiatives are having on their business relative to the resources they have invested into it.

And it doesn't matter how accurate your model is if it's not being widely adopted or not linked to a business goal. Only properly measuring impact can determine the success of your AI initiative.

Key Principles

Measurement

Project success should be determined relative to business outcomes.

You need a clear view on the business impact you need to drive and set your deliverables on this basis. This includes making those delivering the project accountable for impact and not just implementation (aligned incentives protect against tech-focused teams that do not consider the business impact).

The actual measurement metrics will be defined based on the area of implementation. For example, an AI solution embedded within an operational system designed to reduce downtime, or a product that aims to increase customer satisfaction.

Once you've defined your metrics it's important to set up the right performance tracking solution in order you can effectively monitor these over time. This should feed directly back into the data science and machine learning environment to allow for seemless iteration and improvement.

Utilisation/Adoption

You can have the most accurate models and most creative ideas, but if no one is utilising them then they will have zero business impact.

While performance is important, ensure that you put just as much effort into measuring adoption and utilisation rates of AI models across the organisation. This is of course more important in businesses that are inherently built around people and services.

Adoption by usage is one measure, but effective usage of this adoption is more representative of the impact these machines have in aiding and improving human decision making within your business.

Key Questions

1. Do you have best-in-class tools to measure AI impact adopted broadly across the business?

2. Do dedicated roles exist to track the performance of AI within the business?

3. Can you easily get a view of AI metrics across the organisation with dashboards and visualisations?

4. Do you regularly collect data on adoption rates of AI solutions across the business?

5. Are you able to quantify the ROI of your AI investments?

Chapter 4

What Does A Mature AI Strategy Look Like?

A mature organisation can repeatedly take AI ideas from conception through implementation and to generating business value with minimal friction.

By developing your AI culture, ideation, delivery, trust and impact you can start to prove the value of AI to the business and deliver real business outcomes in a way that is iterative and scalable.

Setting A Vision for a Mature AI Strategy

A mature AI strategy is more than just the integration of technology into a business; it's the seamless alignment of culture, ideation, delivery, trust, and impact. It's a journey where each step informs and enriches the next, leading to a future where AI is not just a tool, but a transformative force driving the success of the business.

Our Services

Strategy & Consulting

Align your data and AI ambitions to strategic business objectives and define an actionable strategy that turns data into measurable business value.

Product

Unleash the power of data and AI products for your enterprise. Build differentiated products and services to boost your customer's experience and enhance your operational capabilities.

Data & Software Engineering

Introduce modern data and software engineering practices to pave the way for a scalable operation and extract the true value of data, maximise investment while remaining compliant and secure.

Cloud Platforms

Design, build and deliver data and AI enabled cloud platforms tailored to your organisation's needs, so your focus can remain on adding value to your customers.

Data Mesh

Unlock enterprise data agility at scale with data mesh and develop a highly decentralised data architecture to streamline existing centralised practices to access the data you need, when you need it.

Artificial Intelligence

Go further with your data and discover the benefits from advanced AI. Build upon your data foundations to point AI capabilities at your most pressing business opportunities.

Mesh-AI is a transformation consultancy that exists to reimagine how enterprises operate, making data and AI their competitive advantage.

We turn enterprises into data-driven and AI enabled organisations, unleashing business growth and accelerating outcomes.

Mesh-AI

mesh-ai.com | hello@mesh-ai.com