

LAB³

AI FOR BUSINESS EXECUTIVES

A LAB³ Executive Primer



Digital & App Innovation
Azure

Specialist

Build AI Apps on Microsoft Azure
AI Platform on Microsoft Azure



Microsoft
Solutions Partner

5x Solutions Partner designations
9x Specializations

CONTENTS

Foreword	3
Introduction	4
PART ONE: AI IN GENERAL	
Understanding the Problem	5
Ethical and Governance Considerations	6
Privacy and Sensitive Information	7
PART TWO: TECHNOLOGY	
Why Can't You Just Use ChatGPT	8
Design Tasks: Model Selection	9
Accessing Enterprise Data within Your AI Solution	9
Challenge: Content Gaps and Hallucination	10
Challenge: Knowledge Scattered Across Many Documents	11
Design Tasks: Reasoning	14
Agents Taking Action	15
Multi-Agent Systems	15
Computer Use Agents	16
Bringing It All Together: A Sophisticated Use Case	16
PART THREE: CLOSING THE GEN AI PROJECT	
Testing	18
Sustainment	18
Why Have a Specific 'Landing Zone' for AI?	19
Avoiding Bill Shock	19
Conclusion	20
LAB³ Expertise in AI	21
LAB ³ AI-At-Scale Approach	22
About the Authors	23

FOREWORD

Why Executives need to move beyond the hype and understand AI

By



LACHLAN WHITE
Chief Technology
Officer, LAB³



JASON LEONARD
AI Consulting
Lead, LAB³

While AI has become wildly popular, achieving business value remains a challenge for most organisations with **88% of AI pilots failing to reach production¹**. This has motivated us to put together this primer, to help Executives overcome the challenges with AI and move forward.

To lead AI adoption it is important for Executives to have a high-level understanding of AI developments such that you can make informed business decisions. This will ensure you are well placed to understand the art-of-the-possible and ask the right questions of technologists.

In today's digital era, Executives need to be in a strong position to lead AI adoption across your organisations while avoiding over-inflated expectations.

This primer is aimed at line-of-business executives and senior managers who want a working level understanding of AI capabilities beyond the hype. There are actionable recommendations to help you ensure your organisation gains business value from AI investment.

We hope you find it useful, and please feel welcome to connect with us at any time for further guidance.



Jason Leonard

&



Lachlan White

¹IDC & Lenovo 'CIO Playbook 2025: It's Time For AI-nomics

INTRODUCTION



Millions of people around the world have now experimented with popular Generative AI models like ChatGPT, Claude, DeepSeek, and many others.

Since these models are pre-trained and can be used in a wide variety of circumstances, they have dramatically lowered the cost of AI adoption, putting advanced capabilities within the grasp of businesses of all sizes and even individual employees.

However, the usage of such models to meet corporate requirements typically requires many additional components beyond the model itself, and obtaining an understanding of these components often involves a level of technical detail that is more tailored to technologists than the line-of-business reader.

For example, simply allowing employees to use publicly available models for their daily tasks, exposes organisations to risks, such as leakage of intellectual property or inadvertent release of personally identifiable information (PII); and integrating the AI into business processes and enterprise systems is usually required.

This primer aims to bridge the knowledge divide and empower senior leaders and business executives to confidently make decisions and sign-off on AI investments with a working understanding of the most significant elements of a Generative AI project.

Let's start at the beginning: what are the problems an organisation should prioritise for a potential AI solution?

UNDERSTANDING THE PROBLEM

Albert Einstein is reputed (without evidence!) to have said:

If I had an hour to solve a problem I'd spend 55 minutes thinking about the problem and 5 minutes thinking about solutions.

This is certainly the case in AI projects – with the technology evolving so quickly it frequently grabs headlines and the attention of senior line-of-business executives. In the rush to keep up, jumping to the technology solution often takes precedence over understanding business problems and identifying business benefits.

A well-structured AI project should be the subject of close engagement with business stakeholders to ensure that this is a problem worth solving with AI, and that it is practical and realistic to solve within a given budget and time envelope.

As an executive you should look for a structured approach to this, including bringing stakeholders together in a Business Envisioning Workshop, and using techniques like Design Thinking to thoroughly understand and prioritise the needs of the users before moving to solution definition.

During this solution definition, you will need to understand and factor in non-functional requirements and constraints, such as those imposed by a regulatory body or organisational policies.

Once the business problem is prioritised and clarified, it should also be quantified such that the solution budget can be understood in terms of return-on-investment.

RECOMMENDATIONS

- Start your journey by identifying business challenges that AI might overcome instead of technology PoCs
- Run a Business Envisioning Workshop with LAB³ and Microsoft to ensure you are solving the right problems and have alignment with business stakeholders
- Be clear on the size of business problem, such that the total cost of ownership of a technical solution is aligned.

ETHICAL AND GOVERNANCE CONSIDERATIONS

Along with focussing on the needs of the business and end users, you must also seek to understand constraints – such as privacy, data sovereignty, organisational risk appetite, and more. These should then be extended into discussions with key stakeholders to ensure responsible use of AI, satisfying concerns on fairness, bias, and explainability.

Governments around the world have been grappling with a balance between regulating and encouraging AI development. Using the **Australian Government guidance**² as an example, you should take a risk-based approach and evaluate potential AI implementations on a case-by-case basis, guided by your organisation's AI ethics and governance principles.

Of course, AI Governance does not end when you have decided to green light an AI project; you will need to consider how the AI system will be adopted, used, monitored and governed once in use.

One specific challenge with the use of Large Language Models can be explainability and transparency: just how did the LLM come up with the answer? The internal workings of a LLM are not readily understandable to a human, but fortunately some techniques do exist to provide greater assurance, such as providing citations, or calling out to external deterministic tools (such as calculators or business rules engines).

Let's make this concrete with a common example: an AI-powered customer service chatbot.

For this system to be safe and effective, it must be rigorously tested across various scenarios to ensure it can manage unexpected or complex customer queries. To achieve this, safety protocols must be implemented, including extensive simulation testing, real-world trials, and mechanisms for human oversight.

This system needs to be transparent about its decision-making process. This might involve logging data on how certain responses were generated, which can be reviewed in the event of a customer complaint. Implementing explainability techniques can help understand and trust the AI's decisions and provide opportunities for improvement. This review should be automated, with analytics to systemically check the quality of user interactions, and then used as input for further improvements to the system.

While automated analysis of interactions is a good step, better still is to intervene in problematic exchanges. This is where tools such as Azure Content Safety play a pivotal role with capabilities to detect and block harmful content, such as hate speech, violence, and explicit material, thus avoiding potentially damaging interactions.

A key guardrail is the inclusion of the 'human in the loop' – in the example above, human agents should be able to intervene if needed.

Finally, please keep in mind that ultimate responsibility rests with a person, not an AI algorithm.

RECOMMENDATIONS

- ❑ Publish organisational AI Ethics and Governance principles and create a Chief AI Officer role to own these
- ❑ Name the individual who owns AI governance for each AI application
- ❑ Include 'Responsible AI' as a criterion for project approval.

² [Australia's Artificial Intelligence Ethics Principles | Department of Industry Science and Resources](#)



PRIVACY AND SENSITIVE INFORMATION

A 'ChatGPT for our company' sounds like a good idea to provide general purpose assistance to employees but infers that all content should be available to this new corporate AI tool. This raises risks related to confidentiality and privacy.

Typically, when designing an AI application, you are looking to accelerate a specific business process or increase the productivity of people performing a specific task. These situations do not require access to all corporate data; instead, just the relevant content (for example, just HR Policy information).

On the other hand, since hundreds or even thousands of documents may now be searched more effectively, you need complete confidence that no private or confidential information is contained within them.

Of course, GenAI is not limited to 'a better search' or 'Q&A' use cases. A common use for GenAI is to accelerate the creation of content, including for external publication for marketing or regulatory responses. Business leaders will need to assure themselves that any automated content generation is accurate, explainable, and not inadvertently releasing confidential or private information.

In addition, since the AI application will incorporate new data processing steps and data storage, you need to have guardrails to ensure your data sovereignty and data security policies are satisfied.

RECOMMENDATIONS

- ❑ Limit the AI application to only the data needed
- ❑ Understand how data is managed while it is being processed and stored
- ❑ Understand how sensitive content (like PII) will be detected and decide if redaction or other anonymising techniques are needed
- ❑ Keep track of AI usage in the organisation, such that governance can be applied
- ❑ Keep the 'human in the loop' - humans are ultimately responsible for any generated content, so need to be involved in robust validation.

TECHNOLOGY

Why can't we just use ChatGPT?

Models like ChatGPT, collectively called 'Foundation Models', are trained on a huge amount of information – but not the information that is stored within your enterprise.

As such, the models are great for general purpose questions, but do not have your business content to help with business enquiries that your customers, suppliers, or employees might have about your products and services, status of purchase orders or HR policies.

Moreover, these models are trained at a point in time, so even the answers to general knowledge questions may be out of date.

Further, there is the question of usefulness in a corporate context; while it can be useful to be given some information in a summarised form, you also typically desire some kind of action to take place – a customer might want to extend a warranty; a manager might want to transfer an employee. It is when you move from basic capabilities (such as summarisation, or 'Q&A') to using AI across multiple steps within a business process that you achieve step changes in business performance.

RECOMMENDATION

- To achieve meaningful productivity improvements, target AI to accelerate end-to-end business processes instead of siloed issues.

To engage across a business process, an AI system will need access to corporate data and will need to take action. In turn this will require integration to multiple IT systems in your business.

What follows is an outline of how this can be achieved with additional capabilities and technologies.

For enterprise use, you will also need to take into consideration a range of other topics. These include:

- **Control over the response** provided to the end user, ensuring it is accurate, conveys a tone that is consistent with your branding, and stays on topic
- **Cost monitoring** – as compute is typically charged for on a per-use basis, monitoring of costs is critical to minimising 'bill shock'
- **Guardrails and policies** for cybersecurity, privacy, and intellectual property. For example, how will you prevent leakage of your corporate IP to AI model providers?
- **The need for technologies complementary to AI models** – as you will see throughout this primer, there is a need for architecting and integrating an overall solution to provide a scalable, reliable, and useful outcome to enable adoption.



Finally, you have to use the right tool for the job; as good as these models are, other technologies, including other forms of AI may be a better fit for the task at hand.

Examples that involve significant 'number crunching', such as optimisation, predictive maintenance, and demand and supply prediction are still best to use advanced analytics or machine learning techniques rather than Generative AI.

Design Tasks: Model Selection

Throughout the design of the AI application your team will need to make numerous trade-offs.

As an executive, you will want to make sure your team has adequate experience to recognise these decisions, a specific owner of these technical decisions, and to formalise the considerations behind the decisions.

One of the early solution decisions will be the choice of model to use. The advance in AI Foundation Models has been driven in a large part by an ever larger number of parameters in the model trained on ever greater amounts of knowledge; some of these models are now also 'multi-modal' – i.e. can work with a combination of text, images, and audio. However, these large models also have their drawbacks – the larger the model, the more compute is required, not only for training but also for inferencing (i.e. run-time usage). There is usually extra cost, you can expect increased environmental impact, and potentially slower performance.

Instead, the designer may elect to use a smaller, more targeted model, especially if the pre-trained model knowledge is not as important if the use case is primarily focussed on corporate information and you simply need some natural language processing. As an example, you do not need a state-of-the-art AI model to route emails or understand sentiment in a single language (like English).

RECOMMENDATION

- Ask your team to justify why a large model is being used: smaller models run faster, cost less, can potentially be hosted within your organisation, and have a smaller environmental impact.

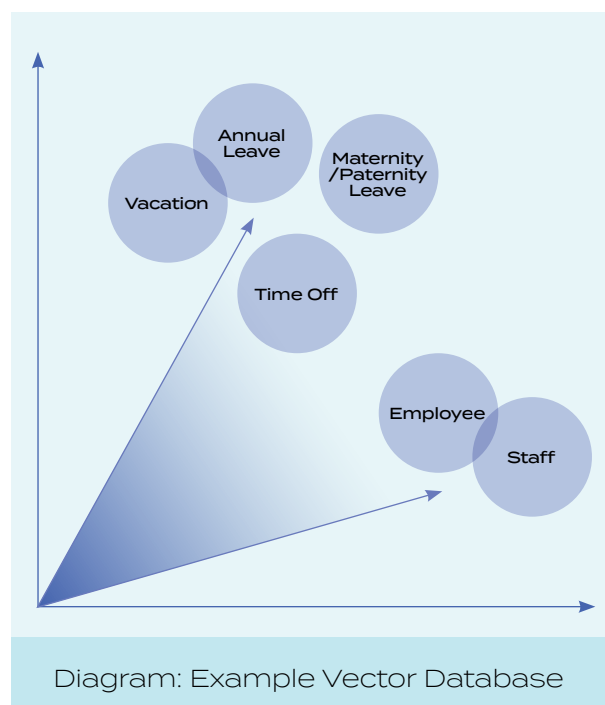
Accessing Enterprise Data within your AI solution

Since you want the AI application to have access to your enterprise content, there will need to be a way for the application to use relevant, up-to-date data. Most corporate data is not stored in structured relational databases, but is in fact unstructured – in text, diagrams, photographs, voice recordings and so on.

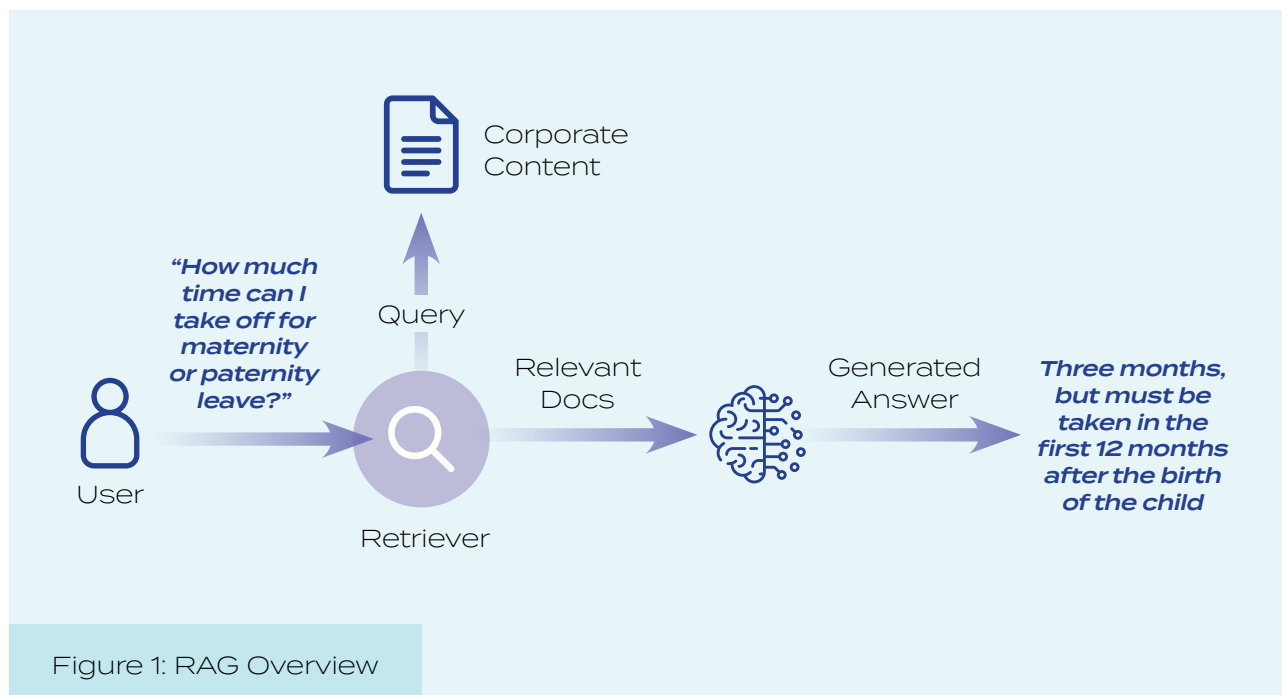
Let's explore a few key technologies for combining your unstructured data with the capabilities of a LLM.

To store unstructured information, the project will use a vector database. Vector databases use a series of numbers ('vectors') to represent parts of text, images etc.

Vector databases are useful to understand similarity; for instance, 'employee' and 'staff' are similar concepts, so vector databases will encode these close to each other. Similarly, 'time off' and 'leave' will also be encoded nearby each other, since the semantic meaning is the same.



In a process called 'RAG' (Retrieval Augmented Generation), your relevant corporate data (e.g. the text in your HR Policies) will be transformed into these vectors and stored in the vector database.



Let's explore how RAG works via an example of an AI-powered HR virtual agent, with a comparison to a traditional search tool.

Managers and employees might be frustrated by the traditional search system, because they have to guess the right key words to get an answer. For example, if they ask about 'leave', but the HR policy documents use the term 'time off', they do not receive any useful results.

In sharp contrast, the AI-powered HR virtual agent can provide an improved experience thanks to the semantic similarity matching: when managers and employees use this virtual agent to ask about the annual leave policy an answer will still be found, even though the original HR policy document only used the term 'time off'.

Challenge: Content Gaps and Hallucination

AI applications acting as expert advisors are only useful when providing useful answers that are otherwise difficult to find.

In practice, when we provide access to end users, they will invariably ask questions that were not contemplated by the authors of the content that the AI has access to! While it can be tempting to blame the AI application, in fact the application is drawing attention to the gaps in content.

To return to the HR Policy example, let's imagine an employee asks for policy information relating to paternity leave. If the organisation has no documented policy for this, then the AI application may provide no answer or even worse, may make it up (known as 'hallucination') if we don't have adequate safeguards in place.

RECOMMENDATION

- ❑ Check availability of Subject Matter Experts (SMEs) to author new content as gaps are discovered.

Challenge: Knowledge Scattered Across Many Documents

Imagine you have content across many documents – for example, customer survey responses or submissions for a local government – and you want to get a sense of the themes involved in the documents.

So far, we have been looking for content that relates to the question – so if there was a document that summarised the themes, the RAG approach would find it, but otherwise it would fall short.

To understand information across multiple sources you can use a knowledge graph. A knowledge graph (stored within a 'graph database') is focussed on entities and the relationships between those entities; an entity could be a company, a person, a place, a product.

Examples of where this is useful include:



Problem Statement	Explanation
Understanding themes from many public submissions <i>(e.g. in government, customer feedback or staff surveys)</i>	AI can analyse and categorise large volumes of text data, identifying common themes and sentiments; a knowledge graph can then link related concepts and entities, providing a comprehensive view of the feedback and highlighting key areas of concern or interest.
Reasoning across multiple documents <i>(e.g. Master Agreements → Statements of Work → Project Change Requests)</i>	AI can use a knowledge graph to connect and contextualise information across different documents. This allows for better understanding of relationships and dependencies, making it easier to track changes, ensure consistency, and identify potential issues or conflicts.
Investigations <i>(e.g. across a large number of safety incidents)</i>	AI can leverage a knowledge graph to correlate data from various sources, uncovering patterns and connections that might not be immediately apparent. This can help in identifying root causes, trends, and potential preventive measures.
Research <i>(e.g. what are all the products our company released in the 1990s?)</i>	AI can query a knowledge graph to retrieve specific information, such as a list of products released in a particular time period. The graph structure allows for efficient searching and filtering, providing accurate and relevant results quickly.
Trends <i>(e.g. across IT support tickets)</i>	AI can use a knowledge graph can link related tickets, categorise them, and link them to other changes, thereby helping to identify root causes.

Let's explore this via an example. Imagine you work in a multinational organisation that provides consultants who have skills in particular industries, and you have the task of finding a banking consultant based in Australia, with past project experience in mortgage processing. Fortunately, you have access to all of the CVs in the organisation, but unfortunately there are way too many consultants for you to remember all the skills and experiences these consultants may have listed in their CVs.

Let's now imagine the prompt for our AI system:

Prompt: "Find consultants with expertise in mortgage processing located in Australia. Highlight those who have significant professional connections and influence within the banking industry. Provide a summary of their key projects and contributions."

A standard RAG pattern might not do too well: recall that it will do similarity matching first, and then an LLM will summarise the resulting short list of CVs. You might end up with three CVs that don't fulfill all the criteria.

A **'Graph RAG' approach³** will instead follow potentially complex relationships, leading to a short list that is more likely to hit the mark. It also has the benefit of being more 'explainable' – an important criterion for use cases where auditability or accuracy is important.

Following is how the relationships might be modelled in a [simplified] graph approach (Figure 2) – an employee works in a country, has past projects in particular knowledge domains, and has expertise in specific industries.

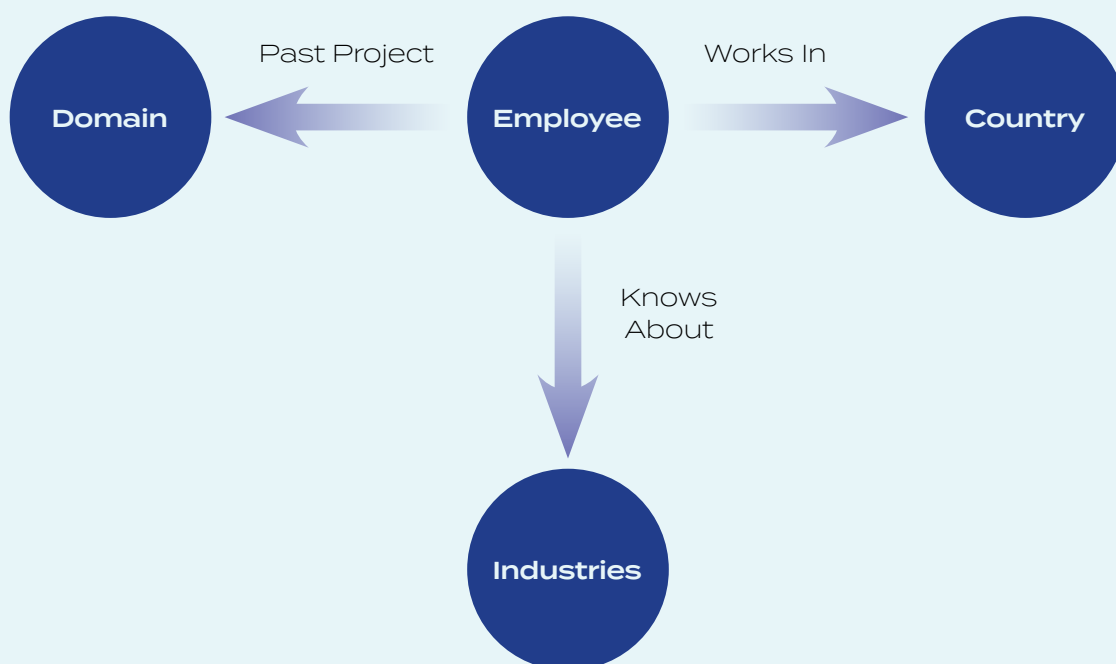


Figure 2: Ontology

³ [GraphRAG: Unlocking LLM discovery on narrative private data - Microsoft Research](#)

When CVs are processed and the results stored in a graph database, part of it could look like what is shown in Figure 3. We see that both Joe and Jon work in Australia, both have experience in banking, but only Joe has past experience in mortgage processing.

As a result, only Joe's CV would be passed to the LLM, which could then summarise his broader experiences and contributions.

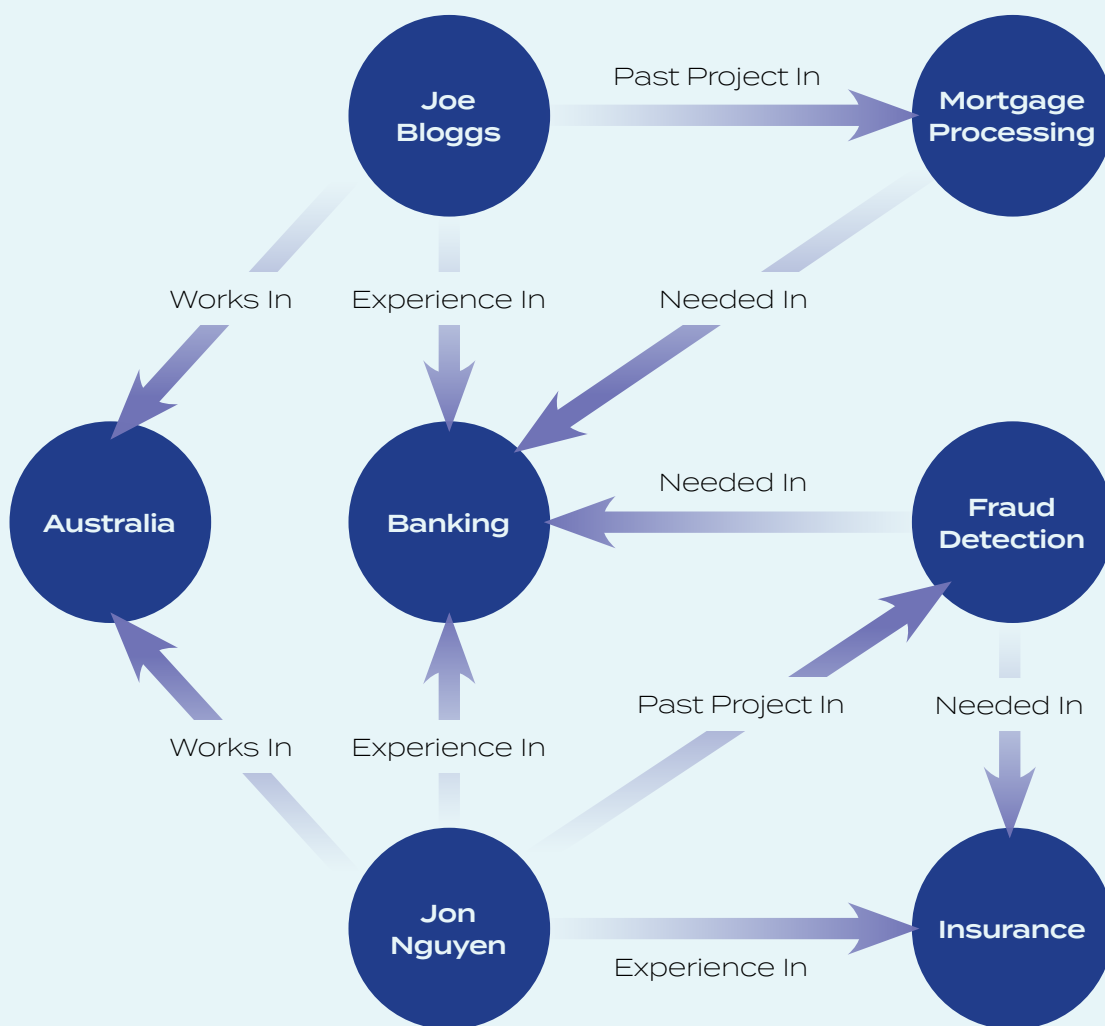


Figure 3: Part of a graph database



Design Tasks: Reasoning

AI is often proposed for time consuming tasks such as making sense of large amounts of content. It is often useful to understand the steps an AI system will need to take to achieve the desired outcome – doing so will help you understand the complexities and feasibility of a project.

For example, we might want an application aimed at helping our HR staff to streamline an understanding of how new government regulations could impact our company policies. In this situation the application will need to:

1. Review external data (the government regulations)
2. Compare to a previous version of the regulations to understand and summarize the changes
3. This summary of changes will then need to be used to find any inconsistencies with your current workplace policies and guidelines.

Contrast this multi-step process with the typical (rather limited) experience interacting with a Large Language Model: you ask a question, you get an answer.

Techniques for a multi-step process include 'chain-of-thought' (CoT) reasoning, Prompt Chaining and Agentic AI.

CoT Reasoning can be implemented via an informed approach to writing the request (or 'prompt').

This can be as simple as telling the LLM to show its reasoning in a step-by-step way or by providing one or more examples.

Prompt Chaining breaks down a difficult request to the LLM to a series of easier tasks. Again, this requires an informed approach to writing prompts.

In fact, a range of sophisticated prompt-based approaches to obtaining a desired outcome have evolved – the takeaway for the executive reader is that developing prompt authoring skills in the organisation can help extract more value from LLMs – some organisations have even formally created the role of a 'Prompt Engineer', and tools are evolving to support this role (e.g. Prompty⁴ for managing prompts).

In short, while AI continues to advance rapidly, use cases that require reasoning are still challenging – after all, AI has tended to be based on statistics and pattern matching rather than following a logic sequence.

RECOMMENDATION

- Ask about how much reasoning an AI use case will need. Does the team have a plan for how to achieve this?

Let's turn our attention to an approach to improving your business outcomes.

⁴ prompty.ai

Agents Taking Action

The usage of an AI application is not always limited to understanding or searching unstructured data such as text and images. In fact, some of the most impactful business outcomes will result from AI supporting an end-to-end business process – which will require the AI to integrate with enterprise systems

Returning to the HR example, let's imagine you want to provide an additional level of service to employees when they wish to book time off – for example, answering personalised questions such as 'How much annual leave do I have?' and taking action, like: 'book 5 days of leave starting on the 12th of December'

AI Agents have the ability to call tools. In this case, the tool would consist of code that knows how to connect to your corporate HR system, extract the leave owing to the employee and potentially book that leave for them.

How does the AI Agent know how to do this? The agent will ask a LLM to understand the nature of the request, and if it matches an action ('book 5 days of leave') then will decide to tell the HR System to do this. It knows which HR System because the agent has this configuration as context. The agent can (and should!) be told to seek user permission before making changes.

Aside from integrating with corporate systems, examples of other actions an AI agent could perform include:

- Web search
- Querying databases
- Triggering a separate automation (e.g. a pre-existing Robotic Process Automation bot)
- Initiating execution of an analytics algorithm.

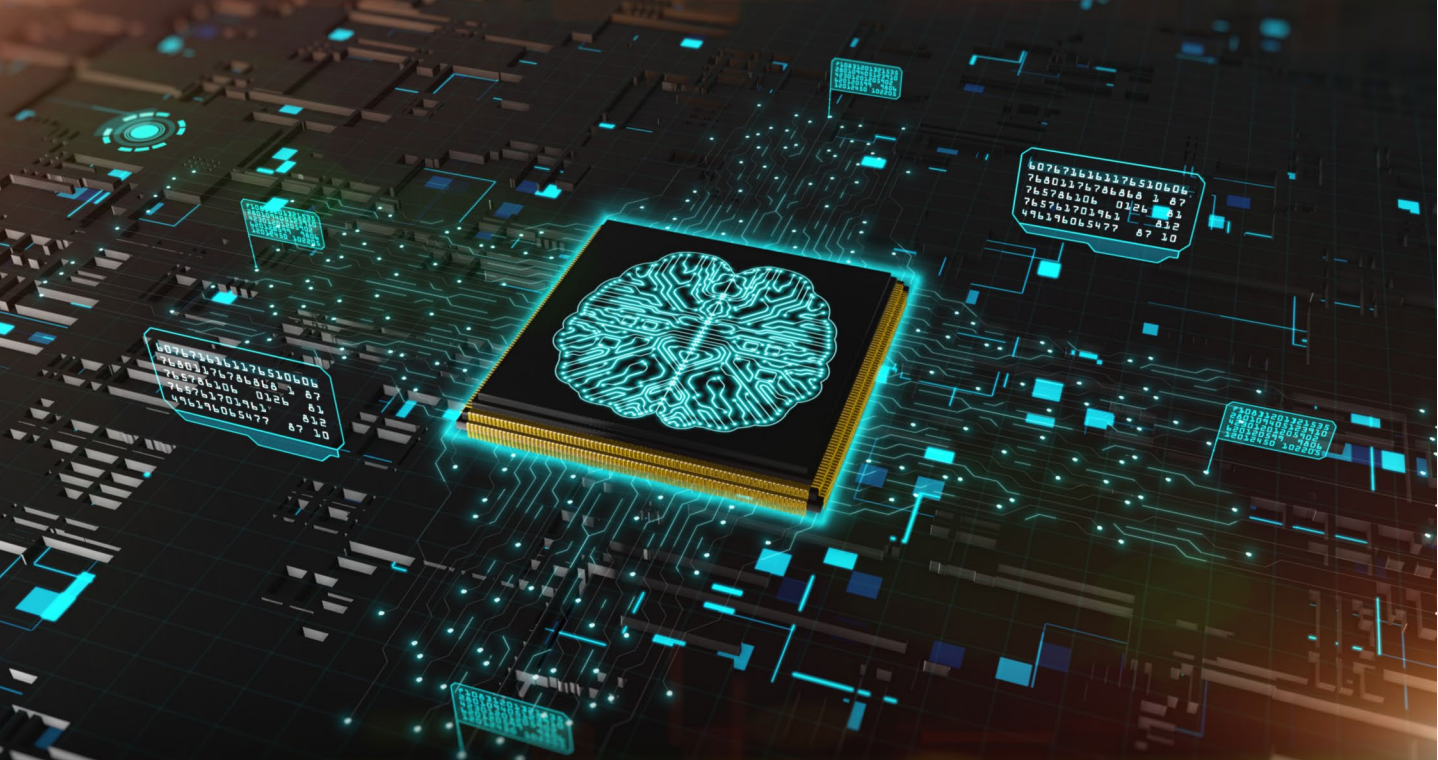
Multi-Agent Systems

A multi-agent approach allows for a complex problem into individual components; once you have individual agents, then further benefits are possible – for example:

- If one agent is performing poorly, it can be improved or replaced as needed
- The agents can use different underlying technologies and can be updated independently as needed
- The interactions between the agents can be tracked, thereby allowing for insights on how a conclusion was reached
- Agents can also keep each other honest – for example, one agent might evaluate the output of another agent. This 'Judge' agent could monitor for hallucination or toxic output.

RECOMMENDATION

- While Agentic AI can be categorised as being 'autonomous' – decide where you need the 'human-in-the-loop' as part of your Responsible AI corporate policy.



Computer Use Agents

Computer Use Agents (CUA) are a new class of rapidly evolving AI agents designed to interact with computer interfaces (including websites) in a way that mimics human users.

They can:

- “See” and interpret screen content
- Use virtual mouse and keyboard actions
- Navigate and interact with applications and websites
- They are able to perform tasks by interacting with the user interface.

CUAs represents a shift towards AI that can directly manipulate and control computer systems through their existing interfaces.

RECOMMENDATION

- If you have a B2C website, check if the design is optimal for navigation by CUAs to avoid missing out on future business⁵.

⁵ See lmstxt.org as an example of one potential change

Bringing it all together: a sophisticated use case

Let's now weave some of these threads together into a more powerful use case.

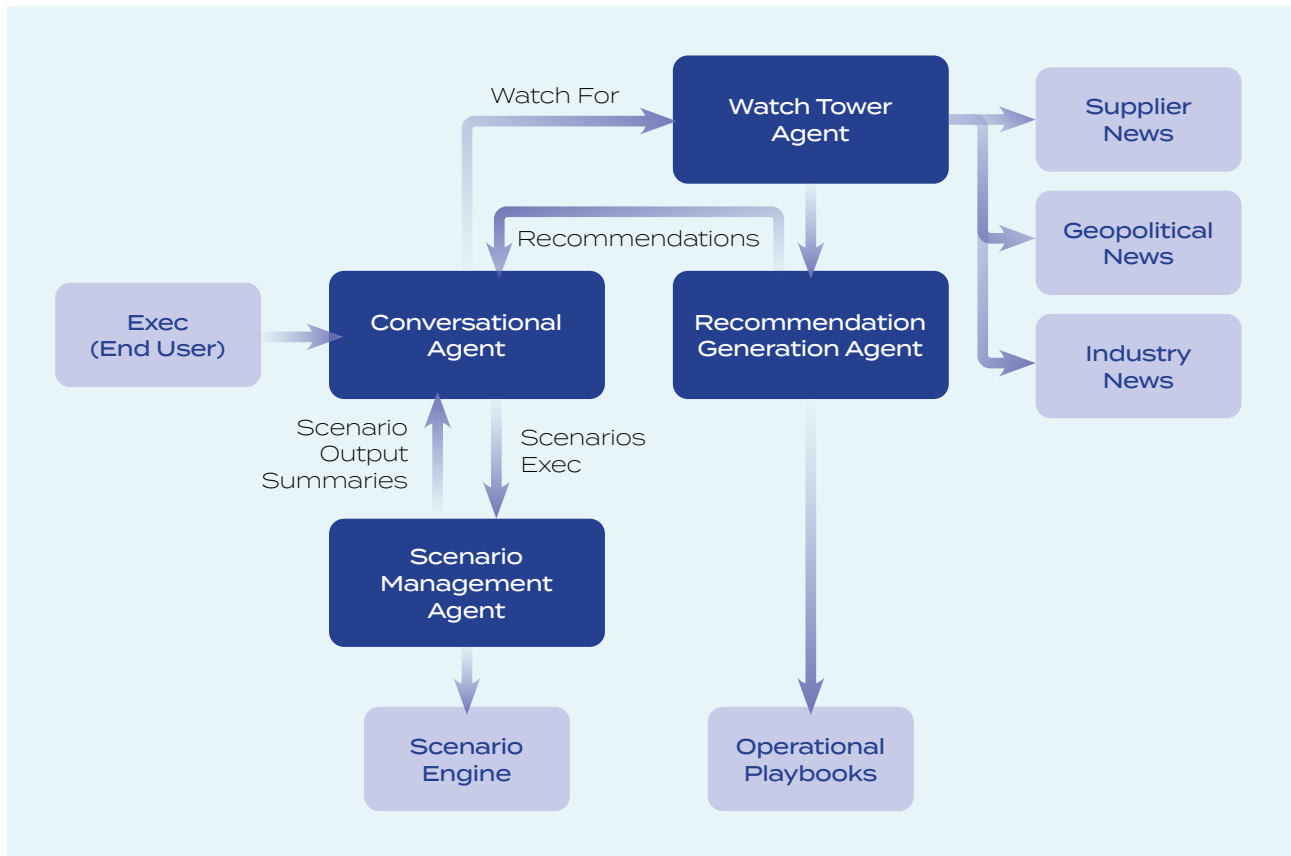
Many knowledge worker roles, including operations centre staff, senior executives and product managers are regularly called upon to make decisions in the absence of complete information – or too much information!

Imagine a scenario where investment decisions need to be made: is demand for our products likely to increase? If so, should the organisation hire more sales reps in a particular city? Or should we instead increase our marketing budget? If we do, will our supply chain keep up with the orders?

We can also imagine the alternative, more difficult scenario: are our products being commoditised or under threat from new market participants?

Ideally an executive would have the support of a tireless advisor, always on the lookout for events and helping the exec to weigh up the options via a conversational interface.

Let's see what such an AI-powered Investment Advisor might consist of:



1. An AI Agent acting as a ‘Watch Tower’.

We ask this agent to look out for market indicators of potential growth or decline in our business. The agent might have access to our internal data (sales figures, customer feedback), public data (social media, product reviews, economic indicators, stock market announcements) and industry sources (industry analyst reports). This agent may combine a LLM with a business rules engine, and a mechanism to alert the executive when necessary. Note this Watch Tower concept might also be applied in an operational context, alerting managers to issues as they arise.

2. An AI Agent to run scenarios.

By tweaking your assumptions, we explore multiple scenarios to explore the likely benefits. These assumptions may be informed by previous decisions (e.g. addition of a sales rep has previously resulted in an extra \$2m per year in sales). Rather than using an LLM, you may use one or more mathematical algorithms.

3. An AI Agent to make suggestions.

These suggestions may be based upon pre-developed playbooks and be triggered by the Watch Tower alerts. It may also trigger scenarios to run.

4. An AI Agent to act as the executive’s conversational interface.

By providing a conversational interface you aim to increase the ability to ‘talk to the data’. This conversational interface can summarize the alert from the Watch Tower, provide the suggestion, represent the likely outcomes as uncovered by the scenario models, and allow the executive to run additional scenarios. Finally, the executive may ask this agent to take a course of action (for example, updating budget allocations in a planning tool).

CLOSING THE GENAI PROJECT



Testing

Traditional IT Systems can be tested on a pass/fail basis. However, Large Language Models are non-deterministic – the same input can lead to a different output!

In addition, an AI solution nearly always requires a level of experimentation to get the best results; and a project team can always try new techniques with a possibility of obtaining even better results, leading to the project continuing indefinitely. This is different to a traditional IT project, which we can consider finished once a set of test cases have successfully executed.

Therefore, a focus on the business outcome and needs of the end users needs to be set at the start and maintained throughout the project, and the test plan focussed on outcomes at this level.

During testing but also in production the team should:

1. Provide a mechanism for end-users to receive feedback
2. Check for hallucination and toxicity
3. Identify out-of-scope usage.

RECOMMENDATION

- ❑ Start the project with the end in mind: what is the definition of success for the AI project?

Sustainment

According to the ancient Greek philosopher Heraclitus “change is the only constant”.

In the world of AI applications, we can expect change in end-user needs, content, models and technologies.

The project to create an initial version of your AI application may have concluded, but to provide long-lasting business value you need a reliable and efficient way to deploy the various components of your application.

You also need to monitor the performance of these applications in production; not only for responsiveness and availability, but also drift, cost monitoring and end-user feedback.

Examples of drift in a customer service scenario:

- Product changes – you can expect questions on the new products
- External events – you can expect questions on whether we are still operating if a major weather event occurs.

RECOMMENDATION

- ❑ Track user interactions with AI applications and take action based on user feedback.



Why have a specific 'Landing Zone' for AI?

Platform considerations can be an afterthought; however, the platform, or 'AI Landing Zone' plays a crucial role in availability, latency, data sovereignty, performance, security and scalability. This platform can then be re-used for each project, thereby significantly reducing the investment and time-to-market for each subsequent AI initiative.

Avoiding Bill Shock

Cost monitoring is important both for avoiding 'bill shock' during production usage, but also potentially for internal accounting purposes (charge-back to the relevant department). Cost monitoring should include tracking inference cost per query. This may lead to a decision to use a lower cost LLM or other cloud-based services.

RECOMMENDATION

- Ensure you select and deploy an AI Landing Zone, such as that provided by LAB³, to reduce project costs and provide security, availability and cost monitoring capabilities.

CONCLUSION

AI projects have a distinct set of technologies, skills and risks. As a business executive, you cannot know all these in detail, but having high-level knowledge can help you avoid some unpleasant surprises.

While this primer contains a number of recommendations for executives, here are a few key take aways:

- Generative AI has lowered barriers to entry for organisations by providing pre-trained models that are useful in a wide variety of situations
- Generative AI systems deal with uncertainty – such as ambiguity inherent in human language. As such, they will make mistakes. Ensure there are guardrails to minimise these mistakes and collaborate with your stakeholders to define what success looks like for your project teams

- Large Language Models, by themselves, are usually insufficient for enterprise needs. Know that you can set your ambitions (reasonably!) higher thanks to frameworks and techniques that have evolved recently, including Agentic AI
- Ongoing production use will need the usual IT and Cloud policies and monitoring – but in addition there are AI specific governance concerns
- Consider forming an AI Centre of Excellence within your organisation to act as the hub for accelerating your AI outcomes by taking ownership of many of the technology platform, business engagement, and governance topics laid out in this primer.

Remember, wherever your organisation might be on your AI journey, Microsoft has recognised LAB³ as a partner with expertise in delivering AI projects for multiple organisations. Don't hesitate to connect with us!



**CONNECT WITH
US AND ENSURE
BUSINESS
VALUE FROM AI**

hello@lab3.com.au



LAB³ EXPERTISE IN AI

LAB³ has been recognised by Microsoft for our AI expertise and our track record in delivering AI projects to multiple clients using Microsoft AI technologies. LAB³ is one of the first partners worldwide to have attained both the Build Apps on Microsoft Azure specialization and the AI Platform on Microsoft Azure specialization.

As you will have seen throughout this primer, bringing together a successful AI initiative needs a combination of:

- **Guided Innovation** – while AI technologies continue to evolve at a dizzying pace, LAB³ makes it our business to connect your business to the right combination of technologies, supported by pre-built technology accelerators and business envisioning workshops to focus on the most impactful initiatives
- **Protection** – robust security throughout the Azure platform allows you to protect your sensitive and confidential information
- **Scalability** – low cost to start, while taking advantage of the global scale of Azure when needed, supported by a pre-configured Azure AI Landing Zone

- **Benefits realisation** – proven ability to move beyond PoCs and help organisations adopt AI at scale and gain measurable benefit from AI.

LAB³ meets organisations wherever you are in your AI Journey, from identifying business value through to use case delivery and establishing an AI Centre of Excellence to scale our business outcomes.



5x Solutions Partner designations
9x Specializations



Digital & App Innovation
Azure

Specialist

Build AI Apps on Microsoft Azure
AI Platform on Microsoft Azure



LAB³ AI-AT-SCALE APPROACH

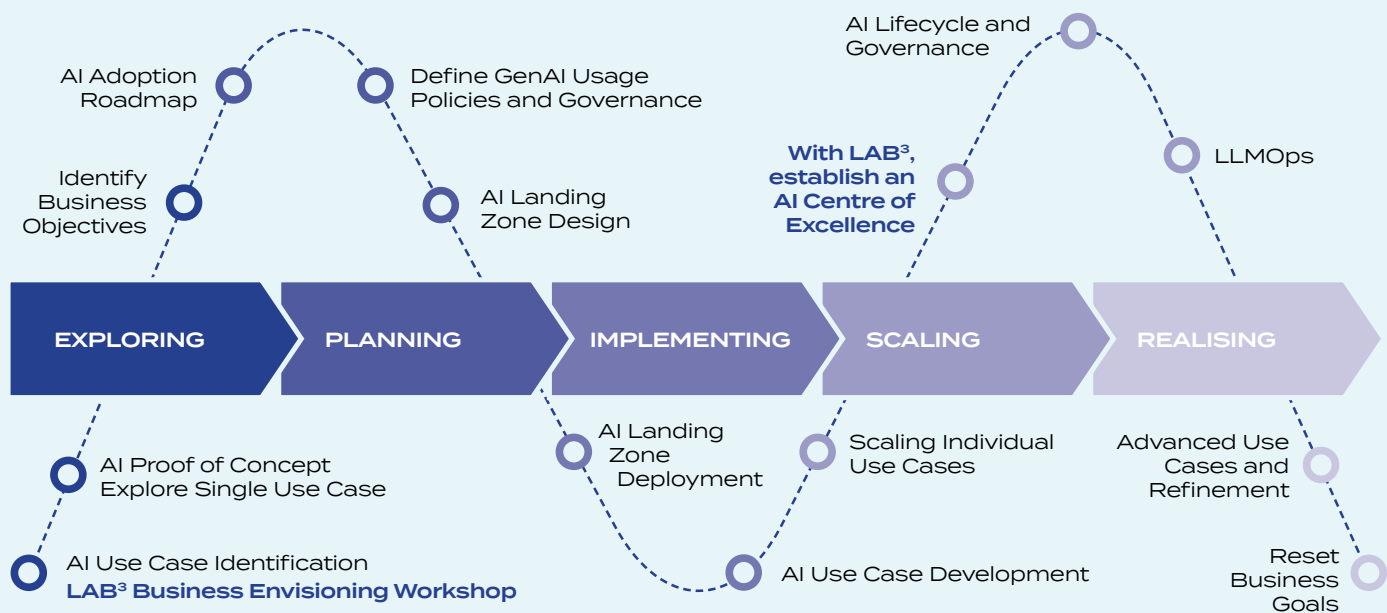


Diagram: Typical AI journey and key actions, with LAB³ providing our services over the whole journey, actioning our AI at Scale framework.

ABOUT THE AUTHORS



JASON LEONARD
AI Consulting Lead,
LAB³

Jason has positioned himself at the forefront of emerging AI technology for over 20 years, helping government departments and enterprises in the Asia Pacific region navigate innovation projects. He is driven to ensure clients achieve measurable business value from scaled AI adoption.

As the AI Consulting Lead for LAB³, Jason is responsible for solutioning, project delivery, market offerings, and for levelling up the skills of our team of data scientists and machine learning experts.

A recent career highlight for Jason includes partnering with an Australian utility to guide an AI at Scale transformation, having oversight of ideation, initiation of all AI, analytics, automation and data projects for asset management. Other recent projects include leveraging generative AI for AI-powered customer service for an insurer, introducing AI to automated coding of pathology in health care, use-cases in Defence, and navigating the complexities of compliance regulations.

With his wealth of insight and experience, Jason is also an Industry Advisory Board Member for RMIT University where he provides advice to the school of Accounting, Information Systems and Supply Chain on industry trends.



LACHLAN WHITE
Chief Technology
Officer, LAB³

Lachlan is a seasoned technologist and a catalyst for accelerated change in large, highly regulated enterprises seeking to modernise their technology platforms and engineering culture. He partners closely with organisations across Government, Financial Services, and Commercial Enterprises in Australia and New Zealand, driving business value through innovative, creative solutions.

As CTO, Lachlan leads LAB³'s technology strategy, accelerators, and R&D. Under his leadership, LAB³ has developed a comprehensive catalogue of unique ready-made accelerators, designed to be deployed across any industry. These solutions build on Microsoft's offerings, enabling clients to accelerate cloud adoption quickly and securely.

Uniquely, Lachlan has been at the forefront of region launches with organisations through all three Australian Azure regions—Sydney, Melbourne, and Canberra. His deep experience has allowed him to guide organisations in maximising value during the release of new regions, helping them leverage these expansions to their strategic advantage. This experience has made him a sought-after expert for organisations looking to take full advantage of the latest cloud innovations in Australia and New Zealand.

Widely respected as an industry expert in cloud and digital transformation, Lachlan specialises in leveraging Microsoft Azure, Terraform, Cloud-Native Technologies, and, more recently, AI. Recognised as a leader in these fields, Lachlan has been a HashiCorp Ambassador for all five years since the program's inception, one of only three worldwide to achieve this. He was also previously an ambassador for the Cloud Native Computing Foundation (CNCF).

Lachlan has also featured as an international speaker, sharing his expertise at various conferences worldwide, where he continues to inspire organisations and technologists in their cloud and digital transformation journeys.

ABOUT LAB³

LAB³ empowers government and business to accelerate, transform and grow fearlessly through controlled, automated, and scalable cloud technology.

Founded in Australia in 2017, LAB³ has experienced hyper-growth and mid 2021 expanded into New Zealand. We are one of the largest Azure engineering practices in Australasia with 250+ skilled employees always pushing over the edge to deliver on any challenge accepted.

With **Microsoft as anchor cloud partner**, LAB³ has a focus on developing repeatable solutions which drive greater value for clients. Significantly, there are world leading ready-made solution frameworks that provide speed and certainty with cloud transformation. At the forefront of innovation, LAB³ ensures the responsible and secure use of Artificial Intelligence to achieve real business value for our clients.

The LAB³ logo is rendered in a white, stylized, handwritten-style font. The 'L' and 'A' are connected, and the 'B' has a distinctive shape with a horizontal bar. The superscript '3' is small and positioned to the upper right of the 'B'.

hello@lab3.com.au

LAB3.COM.AU

©LAB³ Pty Ltd 2025.
All rights reserved.