Queensland Government

FRONTIER SI

SCALE: A Leader's Lens for AI Adoption

A leadership lens for guiding AI initiatives from idea to operations.

Companion Paper 1 to
Governance in the Age of AI: Readiness and Responsible Leadership

Geological Survey of Queensland | FrontierSI
July 2026

Authors:

Mark Gordon (GSQ)
Jia-Urnn Lee (FrontierSI)
Rob Chatterjee (GSQ)
Steven Bowden (GSQ)

Acknowledgements and licence

The Geological Survey of Queensland and FrontierSI respectfully acknowledge the Aboriginal and Torres Strait Islander peoples of Australia, the First Peoples and custodians of the lands, air and waters that sustain the places where we live, work and play. First Peoples have maintained vibrant, living cultures in sustainable relationship with the natural environment for tens of thousands of years, and continue to do so. We pay our respects to Elders past, present and emerging.

This companion paper builds on the original white paper, Governance in the Age of AI: Readiness and Responsible Leadership (GSQ, 2025). The original paper was jointly authored by Dr Jia-Urnn Lee and Gavin Kennedy (FrontierSI), and Mark Gordon, Dr Rob Chatterjee and Steven Bowden (GSQ), with contributions from Professor Flora Salim and Dr Aditya Joshi of the UNSW AI Institute.

The authors would like to both thank and acknowledge the contributions of the following reviewers: Tony Knight, Peter Campbell, and Shoaib Ali from Department of Natural Resources and Mines, Manufacturing, Regional and Rural Development, Graeme Kernich and Annmarie Mabarrack of FrontierSI.

Creative Commons License

The material in this publication is licensed under a Creative Commons CC BY 4.0 – Attribution 4.0 International licence (https://creativecommons.org/licenses/by/4.0), except for third-party material, trademarks, images and photographs. Where a third party holds copyright in this material, copyright remains with that party and their permission may be required for use. Enquiries about this publication can be sent to GSQ via email: gsq@nrmmrrd.qld.gov.au.

More information on this CC BY license is set out at the Creative Commons Website. Enquiries about this publication can be sent to GSQ via email: gsq@nrmmrrd.qld.gov.au.

Citation

GSQ (2026), *SCALE: A Leader's Lens for AI Adoption – Companion Paper 1 to Governance in the Age of AI: Readiness and Responsible Leadership*.

© GSQ 2026

Executive summary

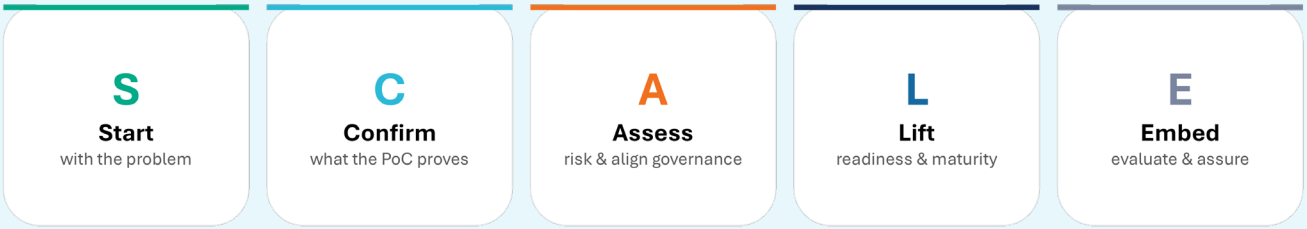

Organisations are experimenting with AI at pace, and many government agencies have now run a proof-of-concept or are planning to.

The hard part is not starting the experimentation process, it is scaling it. Turning a promising experiment into a service the organisation can rely on is where many AI initiatives stall. GSQ met this challenge directly in the delivery of its Digital Librarian proof-of-concept. The experiment worked, but making it into something the organisation could operate, govern and trust has proved a harder and quite different task.

Scaling is where the value is. A proof-of-concept can show that an idea works, and can help organisational change processes associated with AI adoption. However the maximal benefit only arrives when a service reaches production and people come to rely on it. Getting there depends less on the technology than on organisational readiness. The data has to be fit for use, governance has to match the risk, someone has to own the service, and the organisation has to be prepared to run it.

There is already a great deal of helpful and comprehensive guidance on organisational AI adoption in the public domain. SCALE does not compete with any of it. It gives a leader a simple way to engage with this guidance, taking those principles and turning them into a lens that can be used to orient yourself, to frame the conversation across business, technology, governance and risk, and to decide the next step. Additionally, SCALE is not intended to be a replacement for enterprise architecture, cyber security, project delivery and other specialist disciplines. These define the artefacts, designs and implementation pathways required for delivery. The detailed frameworks remain where the depth is. SCALE is the way in.

SCALE is a leadership lens with five elements.

Each element carries a statement that captures its intent. **Start** anchors the work in a real business problem, not in the availability of a technology. **Confirm** recognises that a successful proof-of-concept does not remove complexity, it reveals it. **Assess** keeps governance in proportion, neither over-governing the idea nor under-governing the service. **Lift** treats AI adoption as a question of data, governance, readiness and leadership rather than technology alone. **Embed** closes each stage with a decision, not a deployment, and keeps evaluating the service once it is live.

The five elements stay the same as an initiative moves from idea to experiment, pilot and production, but the answers change at each step, because what is good enough for a proof-of-concept is not good enough for a service people depend on. Used well, the lens helps a leader ask the right questions at the right time, and be honest about when an initiative is ready to go further.

GSQ developed SCALE while scaling its own AI capability, and shares it as a practical starting point for others facing the same challenge. What follows explains each element and shows how GSQ has put it to work.

Contents

Executive summary 3

1 Why this paper 5

2 The SCALE AI leadership lens at a glance 6

Introducing the SCALE elements 7

The origins of SCALE 7

3 The five elements of SCALE 8

S – Start with the problem 8

C – Confirm what the PoC proves 8

A – Assess risk and align governance 9

L – Lift readiness and maturity 9

E – Embed, evaluate and assure 10

4 Bringing it together 11

5 Conclusion 13

Appendix 14

S – Start with the problem 14

What good looks like 14

Supporting guidance 14

SCALE in practice: the GSQ Digital Librarian 15

C – Confirm what the PoC proves 15

What good looks like 15

Supporting guidance 16

SCALE in practice: the GSQ Digital Librarian 16

A – Assess risk and align governance 17

What good looks like 18

Supporting guidance 18

SCALE in practice: the GSQ Digital Librarian 19

L – Lift readiness and maturity 19

What good looks like 19

Supporting guidance 20

SCALE in practice: the GSQ Digital Librarian 21

E – Embed, evaluate and assure 21

What good looks like 23

Supporting guidance 23

SCALE in practice: the GSQ Digital Librarian 24

References 25

1 Why this paper

Everyone is testing AI.
Few are ready to scale it.

Executive leaders in Queensland Government are being asked the same question: "How do you intend to leverage AI in your business?" The pressure to act is high, the technology is moving quickly, and the path to value is often unclear.

Most agencies have already run, or are preparing to run, one or more AI proof-of-concepts (PoC). Only a small proportion of those promising PoCs are reaching pilot and production. This paper is designed to help leaders close that gap.

The barrier to experimenting with AI has fallen sharply as access to AI has democratised. Tools have become inexpensive, accessible and increasingly capable and organisational experience in the application of these tools has lifted markedly. Over the last several years there has been significant work done by governments across Australia to support organisations with AI adoption; however, there are still significant strategic, organisational and data readiness barriers to overcome.

Research reported by CIO found that 88% of observed AI proof-of-concepts did not progress to wide-scale deployment (Schuman, 2025). Separately, McKinsey found that nearly two-thirds of organisations had not yet begun scaling AI across the enterprise (McKinsey & Company, 2025).

.Gartner predicts that, through 2026, organisations will abandon 60% of AI projects that are unsupported by AI-ready data (Gartner, 2025).

AI can highlight where data falls short, but only the organisation can close those gaps. Turning insight into improvement depends on people, process and readiness.

The original GSQ and FrontierSI white paper, Governance in the Age of AI: Readiness and Responsible Leadership (GSQ, 2025), responded to this gap by setting out a number of practical tools, most notably an AI Business Readiness Framework and an AI Capability Delivery Governance Model. These tools were developed to assist GSQ in its organisational readiness journey to scale its Digital Librarian AI proof-of-concept.

Executive leaders do not usually need to begin the journey of AI adoption by implementing a full delivery methodology. Initially they need a way to orient themselves, frame conversations across business, technology, governance and risk, and decide what to do next. This paper introduces SCALE – a five-element leadership lens developed from the lessons and principles explored in Governance in the Age of AI: Readiness and Responsible Leadership and from GSQ's journey in delivering its Digital Librarian AI PoC. SCALE provides a practical pathway for applying those principles across the AI lifecycle, from AI idea and experimentation to accountable delivery.

PoCs create pull.
Readiness creates passage.
Governance creates permission.

GSQ has already completed an initial application of SCALE through its Digital Librarian PoC. That first cycle established technical feasibility, surfaced organisational and governance challenges, and informed a decision to progress towards a pilot. The next cycle of SCALE will be applied during the pilot stage, with different readiness expectations, governance requirements and investment decisions with the goal of achieving a production grade product aimed at a specific audience. In practice, organisations may apply SCALE several times as an AI capability matures from an idea, to PoC, to pilot and to production.

Who this paper is for

- Executive leaders sponsoring, or being asked to sponsor, AI initiatives.
- Senior managers translating AI ambition into a delivery pathway.
- Governance, risk, data and digital leaders working alongside business owners on AI adoption.
- Boards, audit and risk committees seeking a common language to assess AI readiness.

This is a companion paper to Governance in the Age of AI: Readiness and Responsible Leadership. It provides an executive-level view of the journey from AI idea to scalable deployment through the SCALE lens. Where detailed frameworks, criteria or artefacts are referenced, readers are directed to the original paper rather than reproducing the full content here.

2 The SCALE AI leadership lens at a glance

SCALE is a leadership lens. It is not, in its own right, a delivery methodology, governance framework or maturity model. It provides a memorable five-element narrative that helps executive leaders position the work required to move an AI initiative from a successful experiment to a service the organisation can responsibly rely on.

The five elements are presented in a logical sequence – Start, Confirm, Assess, Lift and Embed – to make them memorable. While this sequence reflects the typical progression of an AI initiative, SCALE is not intended to be followed only once. As an initiative evolves, leaders revisit whichever elements are most relevant to each major decision point. For example, they may reassess governance as risks change, confirm the original business case still holds, or identify new organisational capabilities that need to be developed.

Each element therefore represents an enduring leadership principle rather than a single project phase. Together, they provide an assurance cycle that can be applied throughout the AI lifecycle from proof of concept through pilot to production.

Throughout this paper, the AI lifecycle refers to the organisational journey from PoC through pilot to production. While technical AI, software and data lifecycles continue to operate within each stage, SCALE focuses on the leadership decisions required to progress between these stages.

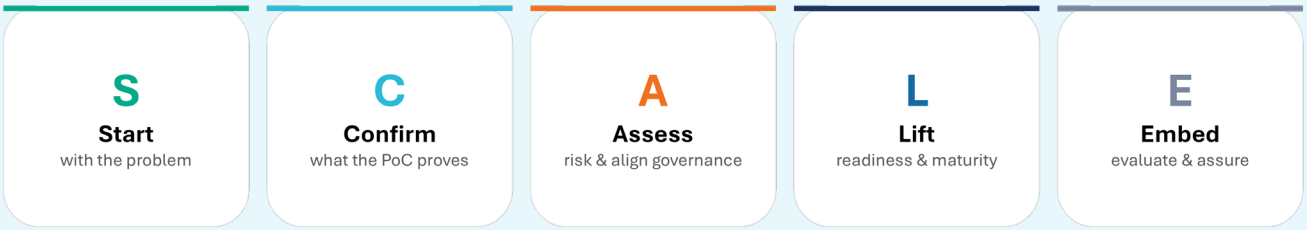

Figure 1: The five SCALE elements – Start, Confirm, Assess, Lift, Embed.

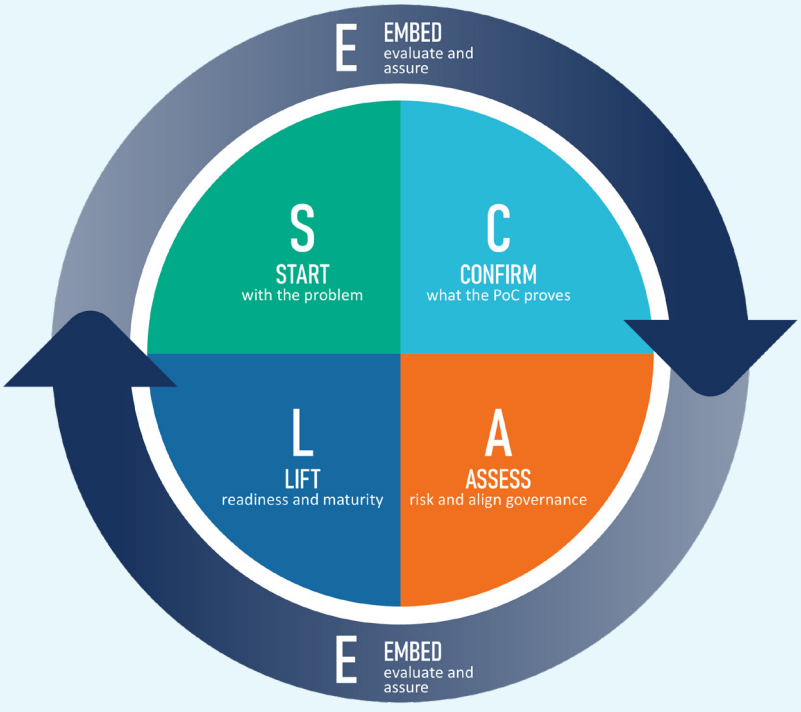

Figure 2: SCALE as an operating cycle. S, C, A and L are revisited continuously through E.

Section 4 shows how the emphasis on each element shifts across the PoC, pilot and production stages of delivery.

Introducing the SCALE elements

Element	What the leader is doing	Element evolution across lifecycle
S – Start with the problem	Anchor the AI initiative in a clear business problem with identifiable users, measurable value and a plausible operational owner.	Regularly revisit the original business problem to ensure the solution remains aligned to the intended value and users.
C – Confirm what the PoC proves	Validate what the proof-of-concept has actually demonstrated, separating technical feasibility from production-grade reliability (measures such as quality, reliability, accuracy of response, customer value and business alignment).	Regularly assess the distance between the PoC, pilot and a production solution to understand uncertainty and risk as the solution matures.
A – Assess risk and align governance	Right-size governance to the level of risk and the current lifecycle stage, avoiding over-governing ideas and under-governing live services.	Review and adapt governance at each lifecycle stage as the risk profile changes.
L – Lift readiness and maturity	Identify and close gaps across Strategy, Organisation, Data and Technology that would otherwise prevent scaling.	Apply the AI readiness review at each lifecycle stage (context dependent) to track progress over time and lift maturity where it matters most.
E – Embed, evaluate and assure	Treat production as an ongoing operating state, with continuous learning, assessment and decision-making across the lifecycle.	Continuously embed, evaluate and assure the service across pilot and production stages, and connect SCALE with the production assurance plan.

The origins of SCALE

While SCALE was developed from GSQ's own experience, each element maps directly onto the frameworks and sections of the original white paper. The provenance of each element is summarised below.

SCALE element	Where it lives in the Governance in the Age of AI: Readiness and Responsible Leadership paper
S – Start with the problem	Section 2.1 Why we did this; Section 2.2.1 Current State (GSQ Open Data Portal)
C – Confirm what the PoC proves	Section 2.2.2 Generative AI Proof-of-Concept; Section 2.2.3 Translating policy into action
A – Assess risk and align governance	Section 4 An AI Governance Framework (Scope 1-4, AI lifecycle alignment)
L – Lift readiness and maturity	Section 3 An AI Readiness Assessment Framework (Strategy / Organisation / Data / Technology domains; spider diagram)
E – Embed, evaluate and assure	Section 5 Summary and Recommendations; Appendix C lifecycle-aligned criteria

If the use case scales, then the organisation must scale with it.

– UQ Executive AI Training Forum, May 2026

3 The five elements of SCALE

This section outlines the what and the why of each S, C, A, L, E element. The full detail – supporting attributes, references and the GSQ worked example – is provided in the appendices.

S – Start with the problem

Anchor the AI initiative in a real business problem, not in the availability of a technology.

The what

The first element anchors the initiative in a real business problem, with identifiable users, a measurable view of the value, and a plausible owner if it succeeds. It also asks the harder question of whether AI is the right approach at all.

The why

This matters because most AI initiatives begin from the technology rather than the problem. If a sponsor cannot describe the business problem in a sentence without mentioning AI, the work is not ready to start. Starting from the problem is what keeps the eventual solution tied to real value and to the people who will use it. AI should be an enabler, not an end in itself.

This element draws on the Digital Transformation Agency's *Guidance for AI Proof-of-Concept to Scale*, which supports defining the problem, aligning the initiative with business priorities, identifying users and intended value, assessing data suitability, considering reuse, and determining whether AI is the appropriate solution. The Queensland Government Digital Investment Principles reinforce the need for initiatives to be achievable, viable and future-focused, while the Digital Service Standard reinforces clear intent and an understanding of users. SCALE brings these considerations together in a leadership lens and adds an explicit human-judgement pathway that is established before the PoC begins rather than after. The leadership lesson is the most straightforward in the paper: start with the problem, not the technology.

Leader's prompt

Before approving an AI PoC, ask: can I describe the business problem in one sentence, without mentioning AI? If not, the PoC is not ready to start.

C – Confirm what the PoC proves

A successful PoC does not remove complexity. It reveals it.

The what

A proof-of-concept answers a narrow question: can this work? Production demands answering a harder one: can we rely on it in a live setting, with real users and real consequences? This element is a deliberate check on what the experiment has actually demonstrated, separating technical feasibility from production-grade reliability.

The why

A successful PoC does not remove complexity; it reveals it. A short coherence check before the next round of investment is one of the highest-value things a leader can ask for. A coherence check is a structured review that tests whether the PoC actually holds together as a basis for further investment – whether it solves a real problem, whether users trust it, whether the data and governance will support production, and whether anyone is ready to own it.

The coherence check sits alongside established investment and delivery disciplines, including the Queensland Business Case Development Framework, the Treasury Project Assessment Framework (PAF), and methodologies such as PRINCE2, together with the post-PoC section of the DTA guidance. SCALE adds a single prompt to these – risk and reliance, in particular whether users could come to over-rely on the AI outputs. The leadership lesson is to treat a successful PoC as the start of the real questions, not the end of them.

Leader's prompt

If the answer to most of these prompts is "we'll work that out in production", the initiative is not ready to leave the PoC stage. It is ready for a more deliberate next step – usually a pilot, with the next layer of governance and readiness in place.

A – Assess risk and align governance

Don't over-govern the idea. Don't under-govern the service.

The what

Governance is the point at which good initiatives most often stall, usually because it is applied as a fixed burden rather than scaled to the actual risk and stage of the work. The A element of SCALE right-sizes governance to risk, enough to be safe and accountable, without so much process that early exploration grinds to a halt.

The why

There are two failure modes to manage at the same time. Over-govern the idea and teams stop bringing ideas forward, or route around the process. Under-govern the service and a successful experiment drifts into real use without the accountability that production requires. Risk and exposure should decide when each layer of governance becomes essential.

This element is supported by the Queensland FAIRA framework as the entry point for risk assessment, the DTA Policy for the Responsible Use of AI in Government, and the National Framework for the Assurance of AI in Government (Australian Government et al., 2024), with ISO/IEC 42001, 42005 and 38507, the Privacy Act and the relevant security frameworks applying as risk and exposure grow. SCALE stays deliberately agnostic about which risk method you use. The leadership lesson is to let risk drive governance, not the other way around: do not over-govern the idea, and do not under-govern the service.

Leader's prompt

For your initiative right now: what level of risk and organisational impact is it operating at? What level of risk and organisational impact are you investing for? If those two answers do not match, governance and investment are misaligned.

L – Lift readiness and maturity

AI adoption quickly becomes a data, governance, readiness and leadership problem.

The what

AI adoption should quickly stop being a technical problem and become a question of data, governance, capability and leadership. This element identifies the gaps across Strategy, Organisation, Data and Technology that would otherwise prevent an initiative from scaling, and is honest about which gaps have to close before the next stage.

The why

Technical capability and organisational capability rarely mature at the same rate, and strength in one area can hide a critical gap in another. This is the point at which organisations often draw on enterprise architecture practices to align strategy, organisation, data and technology into a coherent operating capability. This is also where the difference between being data-ready and being AI-data-ready tends to surface, often as the thing that decides whether a promising PoC can actually scale.

This element builds on the AI Readiness Assessment Framework in the original white paper, which scores Strategy, Organisation, Data and Technology across four readiness areas, and is complemented by the AI Plan for the Australian Public Service, the Australian and Queensland Government Architectures, and the DTA readiness guidance. The leadership lesson is that AI adoption is a multi-domain problem: the four domains must scale together, because strength in one can mask a critical gap in another.

Leader's prompt

Ask the people accountable for Strategy, Organisation, Data, Technology to provide an honest current-state score and a target-state score for the next lifecycle stage. These may be executive leaders, business owners or functional leads depending on your organisation. The conversation that produces those numbers is usually more valuable than the numbers themselves.

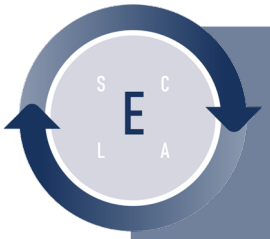

E – Embed, evaluate and assure

Every lifecycle stage ends with a decision, not a deployment.

The what

A service is not deployed once and left alone. It operates in a setting where users, data, models, regulation and risk all keep changing. The ‘E’ element proposes continuous evaluation and an explicit decision at each lifecycle stage: proceed, pause, redesign, scale or stop.

The why

The point for a leader is that every stage should end with a decision, not just a deployment. SCALE is meant to be run again as reliance and risk grow, so the governance and readiness that suited a PoC are revisited before the service is trusted in production. The goal of AI deployment is an accountable operational capability, rather than a focused technology product alone. Achieving this relies on continuous measurement, and on applying the results of that measurement to course-correct as risk and readiness change. At each stage of the delivery lifecycle, leaders must evaluate whether the initiative continues to deliver value, and assure themselves that governance, risk, data quality and operational controls remain fit for purpose.

This element aligns with the DTA Technical Standard for Government’s Use of AI and its Discover, Operate and Retire lifecycle, the NIST AI Risk Management Framework, ISO/IEC 42001 and ISO/IEC 5338, and ITIL continual improvement.

The leadership lesson is that the end product is an operating state, not a destination. Build the evaluation loop and use it at every stage.

Leader’s prompt

For a service already in production, ask: when did you last make an explicit decision to proceed, pause, redesign, scale or stop? If it is simply running, evaluation has quietly stopped.

4 Bringing it together

SCALE is intended to give leaders a practical entry point into the methodology set out in Governance in the Age of AI: Readiness and Responsible Leadership. It is not a replacement for that methodology. The strength of SCALE is that it is memorable and easily communicated; the strength of the underlying white paper is that it provides a framework that complements existing guidance, policies and standards to support AI adoption. Leaders should apply all of these things together.

SCALE does not prescribe how organisations integrate the Strategy, Organisation, Data and Technology readiness domains. Many organisations already do this through enterprise architecture, while others achieve it through existing governance, portfolio or operating model practices. What is important is that these domains are considered together as the organisation matures an AI capability.

Each stage of AI capability development creates new questions, new risks and new readiness requirements. What determines success at the proof-of-concept stage may be quite different from what determines success at pilot or production. SCALE is therefore intended to be completed multiple times, revisited whenever organisational reliance, risk, investment or maturity changes.

SCALE is not strictly a linear process or roadmap, but it does have some elements of order when applied to delivery. An important part of the conversation is understanding the level of investment or effort required at each stage. The following figure unpacks these considerations in the context of the proposed delivery of the GSQ Digital Librarian capability, across three stages: PoC development; scaling the PoC to a pilot capability with a reduced stakeholder group; and the delivery of a full production instance as part of an accountable operating model.

Applying SCALE across the AI lifecycle

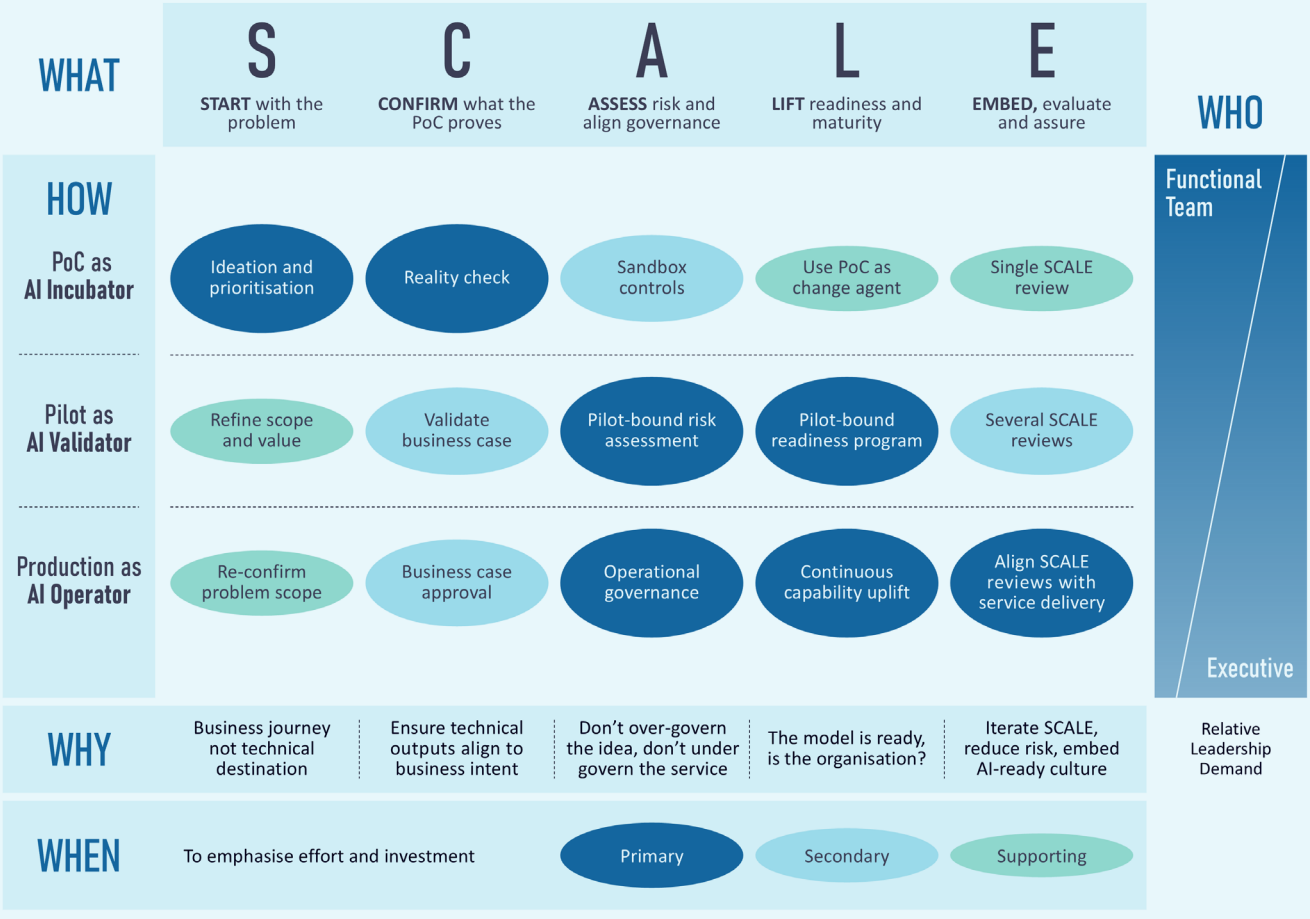

Figure 3: Relative organisational investment and effort on each SCALE element across PoC, pilot and production. Effort is cumulative.

At each PoC-Pilot-Production stage, the level of investment or effort in each S-C-A-L-E element varies with the emphasis and commitment required, and effort is cumulative as the tools are applied through delivery. For a PoC or experiment, the majority of investment is in the S and C elements initially, with the pilot and production stages adjusting and confirming the problem statement and business case to incorporate what was learnt at the previous stage.

The importance of the A, L and E elements increases in the pilot and production stages, as risk increases, governance is scaled to meet that risk, and the organisation is prepared for deployment with tighter and more frequent assessment cycles to ensure delivery is assured appropriately.

It is also helpful to think of the PoC stage as an incubator, the pilot stage as a validator, and the production stage as the operator stage. Individual applications can vary the approach – the application of SCALE is genuinely contextual. Some initiatives may not use a piloting stage at all, and the frequency of risk, governance and readiness assessments may vary. It is up to the leader to decide what is appropriate, considering risk and organisational maturity.

Leadership lessons

1. Start with the problem, not the technology. If you cannot describe the business problem without mentioning AI, the initiative is not ready to start.
2. Treat a successful PoC as the beginning of the questions, not the end. A coherence check before the next stage of investment is one of the highest-value activities a leader can sponsor.
3. Don't over-govern the idea, and don't under-govern the service. Risk-proportionate, lifecycle-aligned governance is the difference between innovation that scales and innovation that stalls.
4. AI adoption is a multi-domain problem. Strategy, Organisation, Data and Technology all must scale together. Strength in one domain can mask a critical gap in another.
5. Production is an operating state, not a destination. Embed, evaluate and assure the service – and close every lifecycle stage with an explicit decision.

5 Conclusion

Most leaders reading this are already being asked what they are doing with AI. The easy response is to point to a successful experiment. The more valuable one is to show a credible path from that experiment to something the organisation can rely on, and stand behind, in live use.

SCALE does not make that path shorter. What it offers is a way to see the whole path clearly, and a shared language for working through it with delivery teams, governance bodies and the executives who sponsor the investment. Each S-C-A-L-E element is simple enough to hold in mind, and each one points to the detailed frameworks in the appendices and the original white paper when the work gets specific. Start with the problem, not the technology. Confirm what the experiment has actually proven before treating it as permission to scale. Assess risk and let it drive governance, rather than applying governance as a fixed burden. Lift readiness across all four domains of Strategy, Organisation, Technology and Data, because strength in one cannot compensate for a gap in another. And embed, evaluate and assure the service as an operating state, closing each stage with an explicit decision rather than a silent drift into reliance.

The most important shift SCALE asks of a leader is in how the destination is understood. Is it a technology milestone reached once and left behind, or is it an evolving state that has to be earned, and re-earned, as users, data, risk and organisational reliance change. That is why SCALE is iterative. A PoC is assessed against the question “can this work”? A pilot is assessed against “can this be operated”? A production service is assessed against “can this be sustained and scaled responsibly”? The same five elements apply at each stage, but the readiness expectations, governance obligations and investment decisions are different each time. Applying SCALE well means knowing which questions matter most at the stage you are actually at, and being honest when the answer is that you are not ready to progress.

GSQ is on this path with the Digital Librarian. The first SCALE cycle established technical feasibility, surfaced the organisational and governance challenges that a production service would have to resolve, and ended with an evidence-based decision to prepare a pilot rather than rush to production. The second cycle is now beginning. We are sharing SCALE not as a finished answer but as a practical, portable way of thinking that other agencies may adapt to their own context. As a lens, it is deliberately framework-agnostic: it sits above the specific risk, investment and assurance frameworks each organisation already uses, and connects them into a single narrative a leader can communicate.

This is the first in a planned series of companion papers building on Governance in the Age of AI: Readiness and Responsible Leadership. Future papers will go deeper on enterprise AI capability, the strategic timing of investment, managing change load across Run-Grow-Transform, and information stewardship as the foundation of responsible AI. SCALE is the entry point; the work it points to is where the real value is realised.

GSQ welcomes feedback on the value other organisations find in the approach, and on how SCALE can be sharpened as our own journey with the Digital Librarian continues.

Appendix

The appendix expands on each SCALE element by providing supporting attributes, references, and a narrative drawn from GSQ's Digital Librarian initiative. These are intended as supplementary material for leaders and delivery teams who want to apply SCALE in more detail.

S – Start with the problem

Anchor the AI initiative in a real business problem, not in the availability of a technology.

AI adoption is a business journey rather than a technical destination, and so any AI use case needs to be business-led. The application of SCALE begins by returning the conversation to the first question every executive should ask: what specific business problem is this AI initiative designed to solve, and is AI actually the right approach?

What good looks like

A good AI PoC problem will usually have most of the following attributes.

Problem and value

- Clear business problem. A specific workflow, decision, service or pain point – not a vague desire to "use AI".
- Strategic alignment. Maps to agency priorities, service improvement, productivity, risk reduction or public value.
- Measurable value hypothesis. The value proposition is clear: time saved, quality improved, risk reduced or decisions better supported.

Users and operations

- Identifiable users. Clear users, beneficiaries and decision-makers can be engaged.
- Plausible operational pathway. There is at least a candidate owner, workflow and support model if the PoC succeeds.
- Service vs process distinction. Are we using AI to improve a service, or just to speed up a flawed process?

Data and risk

- Data availability and suitability. Relevant data exists, is accessible enough to test, and its quality, provenance and sensitivity are understood for the stage.
- Proportionate risk. The initial test can be designed safely, with bounded exposure to users, data and downstream impact.
- Human judgement pathway. The role of expert review, validation or decision-making is clear from the start.
- Appropriateness. Does the problem statement create any ethics, privacy, confidentiality, inclusion, offence or legal issues?

Potential to scale

- Repeatability or leverage. The learning, pattern or capability could support other use cases, not just this one.
- Process novelty check. Has this been done elsewhere, and if so, can we reuse rather than rebuild?

Supporting guidance

Several of these attributes draw on Appendix 1 of the DTA's *Guidance for AI Proof-of-Concept to Scale* (March 2026), including a clear business problem, strategic alignment, identifiable users, measurable value, data availability and suitability, reuse, and consideration of whether AI is the appropriate solution. Additional supporting principles are provided by the Queensland Government Digital Investment Principles, including that investments should be achievable, viable and future-focused. Appendix 2 of the DTA guidance provides examples of common AI use cases and implementation scenarios.

SCALE distils these concepts into a leadership-focused set of attributes for ideating and prioritising AI initiatives. It also makes the human-judgement pathway an explicit consideration at the problem-selection stage: the role of expert review, validation or decision-making should be established before a PoC begins rather than after. Not every attribute needs to be fully developed during early experimentation, but most should be clearly present before the PoC is approved.

Element	Primary (external)	Additional references (external)
S – Start with the problem	DTA (2026) – *Guidance for AI Proof-of-Concept to Scale*; Queensland Government – Digital Investment Principles	DTA – Digital Service Standard (clear intent; know your user)

SCALE in practice: the GSQ Digital Librarian

Problem and value

GSQ provides access to approximately 191,000 datasets through the GSQ Open Data Portal, including more than 105,000 report-category records. Across its data platforms, GSQ hosts more than 500 terabytes of geological data. These reports are rich, lengthy and often the only source of specific insights. A single user question – for example, "provide me with an exploration history in this area" – can require manual review of tens or hundreds of reports. This challenge aligned directly with GSQ's role as the custodian of Queensland's geological information. GSQ sought to improve the discoverability and use of existing public geoscience data by using generative AI to improve conformance with the FAIR data principles (Findable, Accessible, Interoperable, Reusable) (Wilkinson et al., 2016).

Users and operations

The Digital Librarian PoC was not framed as "let's try GenAI". It was framed as a knowledge-discovery problem: how do we enable users to obtain rapid, traceable insight from a very large unstructured corpus, while retaining the human expert in the loop? The intended users are resource explorers, geoscientists and other stakeholders seeking to navigate large volumes of geological reports, find and synthesise information faster, and focus more on analysis and decision-making. The Digital Librarian was designed to support expert judgement, with users remaining responsible for validating AI-generated responses against the original reports.

Data and risk

GSQ entered the Digital Librarian PoC from a position of relative strength in data readiness. Years of investment through the GSQ Data Modernisation program had established the Open Data Portal, metadata standards, data governance processes and trusted custodianship arrangements. This foundation made it possible to explore AI-enabled data discovery without first solving fundamental data access and quality challenges. Risk was managed by constraining the PoC to a controlled environment using a curated data repository (sandbox) and internal users.

Potential to scale

The Digital Librarian was selected because success would create value beyond a single PoC. If successful, it could transform how users discover, interpret and apply geological information across Queensland's resource sector. As the challenge already existed within GSQ's core information services, there was also a plausible pathway for operational ownership if the capability proved valuable. This framing made the Digital Librarian a credible candidate for AI. It also provided a reusable template for future AI initiatives, demonstrating how data retrieval, governance and expert review could be combined for other information-intensive use cases across government.

C – Confirm what the PoC proves

A successful PoC does not remove complexity. It reveals it.

PoCs are designed to answer a narrow question: can this actually work? Production demands answering a harder question: can we safely rely on this in a live operating context? The transition between these two questions is where many AI initiatives quietly stall.

Many organisations interpret a successful PoC as evidence that they are ready to proceed directly to production. Realistically, a PoC often reveals the next set of questions that need answering before larger investments can be justified. This is in part why SCALE is an iterative process, with regular cycles to revisit delivery and ensure that scope, risk, governance and readiness parameters are optimised for a successful outcome. These questions should be explored through an appropriately sized business case development process, as part of an endorsed delivery methodology such as PRINCE2 (PeopleCert International Limited, 2023).

What good looks like

Before approving the next stage of investment, leaders should put any successful PoC through a coherence check. Outlined below is a starter set of questions drawn from GSQ's experience, the original white paper (Appendix C) and the reference list. The nature of any additional questions should be tied to the intent and context of the proposed delivery.

Value and use-case fit

- Use-case fit. Does the PoC solve a real business problem? Is AI the right approach, or is the underlying issue search, workflow, metadata or data quality?
- User validation. Do real users find it useful, usable and trustworthy? Does it change how they work or make decisions?
- Benefit evidence. Is there evidence of time saved, quality improved, confidence increased or decisions better supported?

Performance and trust

- Answer quality. Are outputs accurate enough? Are citations, sources and reasoning pathways visible enough to support confidence?
- Data performance. Is the source data fit for purpose? Are gaps, bias, duplication, poor metadata or provenance issues affecting results?
- Risk and reliance. What happens if the answer is wrong, incomplete or misunderstood? Could users over-rely on it? What is the real impact?

Scale and sustainability

- Governance readiness. What privacy, cyber, legal, records, ethical or policy obligations are triggered if the use case moves beyond PoC?
- Operating feasibility. Who would own, fund, support, monitor and improve the service if it moved to pilot or production?

Supporting guidance

The PoC coherence check aligns with broader Australian Government and Queensland investment guidance. The post-PoC section of Appendix 1 of the DTA's *Guidance for AI Proof-of-Concept to Scale* supports assessing whether the initiative continues to address a real problem, demonstrates value, has been validated with users, remains technically feasible, and has a credible pathway towards implementation. Queensland's Business Case Development Framework complements this guidance through its staged approach to investment analysis, benefits, governance and decision-making.

SCALE adds one prompt to these: risk and reliance – in particular, whether users could come to over-rely on the AI's outputs, which we consider essential to test before next-stage investment is committed.

Element	Primary (external)	Additional references (external)
C – Confirm what the PoC proves	Queensland Government – Business Case Development Framework (Stages 1–3); Queensland Treasury – Project Assessment Framework; DTA – PoC-to-Scale Guidance, Appendix 3 (procurement);	AXELOS (2023) – PRINCE2® 7

SCALE in practice: the GSQ Digital Librarian

Value and use-case fit

The Digital Librarian PoC confirmed that generative AI could help users interrogate large collections of unstructured geological reports and retrieve useful information more efficiently. It demonstrated that AI was an appropriate approach to the business problem originally identified.

Performance and trust

The PoC also clarified what it had not proven. It had been deliberately scoped within bounded technology, data, personnel and processes (GSQ, 2025, Sections 2.2.2 and 4). Moving beyond those bounds raises a different class of questions about data quality and provenance, accuracy and reliability, governance, and enterprise integration. These questions will become the focus of GSQ's next SCALE review.

Scale and sustainability

Perhaps the most significant outcome was that the emergent challenges were organisational, rather than about improvements to the AI model. Scaling the Digital Librarian would depend on data stewardship, service ownership, and how an AI-enabled tool would integrate with existing Queensland Government enterprise and corporate governance frameworks (GSQ, 2025, Section 2.2.2). These questions quickly became more significant than questions about the underlying AI model. The PoC therefore became a catalyst for broader organisational change, and informed the development of the readiness and governance approaches presented in the Governance in the Age of AI paper.

A – Assess risk and align governance

Don't over-govern the idea. Don't under-govern the service.

Governance is often the point at which successful AI PoCs stall. This is not because governance is unnecessary, but because it can become an overwhelming burden when it is not calibrated to the initiative's risk profile, maturity and stage of delivery. In many cases, leaders and delivery teams are also unclear about which governance obligations apply, when they should be introduced, and who is responsible for completing them. The original white paper addresses this challenge by proposing a risk-aligned governance model that helps leaders understand their obligations, sequence governance activities across the delivery lifecycle, and assign clear accountability for the work required.

The two governance traps

Leaders who strike the right balance manage two opposing failure modes at the same time.

Trap	What happens	Consequence
Over-govern the idea	Early exploration is treated like a production system, with the full weight of enterprise governance applied at the outset.	Innovation slows or stops. Teams stop bringing forward ideas, or route around governance.
Under-govern the service	A successful PoC is allowed to drift into real use without the governance, accountability and support model that production requires.	Risk grows without clear accountability. Trust, compliance and reputation are exposed.

Controls should grow with risk, maturity, exposure and investment as the initiative progresses through the lifecycle.

Right-sizing governance to risk and scope

The original white paper provides a detailed, tiered AI Capability Delivery Governance Model that maps governance artefacts to four scope tiers—from very low to high risk—and to the six lifecycle stages used in the original white paper (GSQ, 2025, Section 4 and Figure 5).

By applying a risk assessment of choice (for example, FAIRA or NIST), leaders can quickly establish which governance artefacts are most important to the risk level they have defined, when to implement that governance, and who is responsible for delivering those outputs.

Scope tier	Example	Governance posture
Scope 1 – Very low	A commercial GenAI tool used on publicly available open data.	Minimal: an accountability framework, and a risk and impact management plan.
Scope 2 – Low	An internal PoC in a firewalled environment with curated data and internal users only (e.g. the Digital Librarian PoC).	Adds a regulatory and legal plan, and a privacy protection and security plan.
Scope 3 – Moderate	A productionised text-based service with beta access for trusted external users.	Adds an AI assurance plan, a fairness plan, and a product strategy / management plan.
Scope 4 – High	A self-trained, multi-modal, public-facing service trained on all available data.	Adds a contestability and human oversight plan with robust human-in-the-loop protocols.

At each tier, all governance artefacts from the lower tiers remain required (GSQ, 2025, Section 4.4). The model is cumulative, not optional. The point is that risk and scope determine when each layer becomes essential – not the calendar.

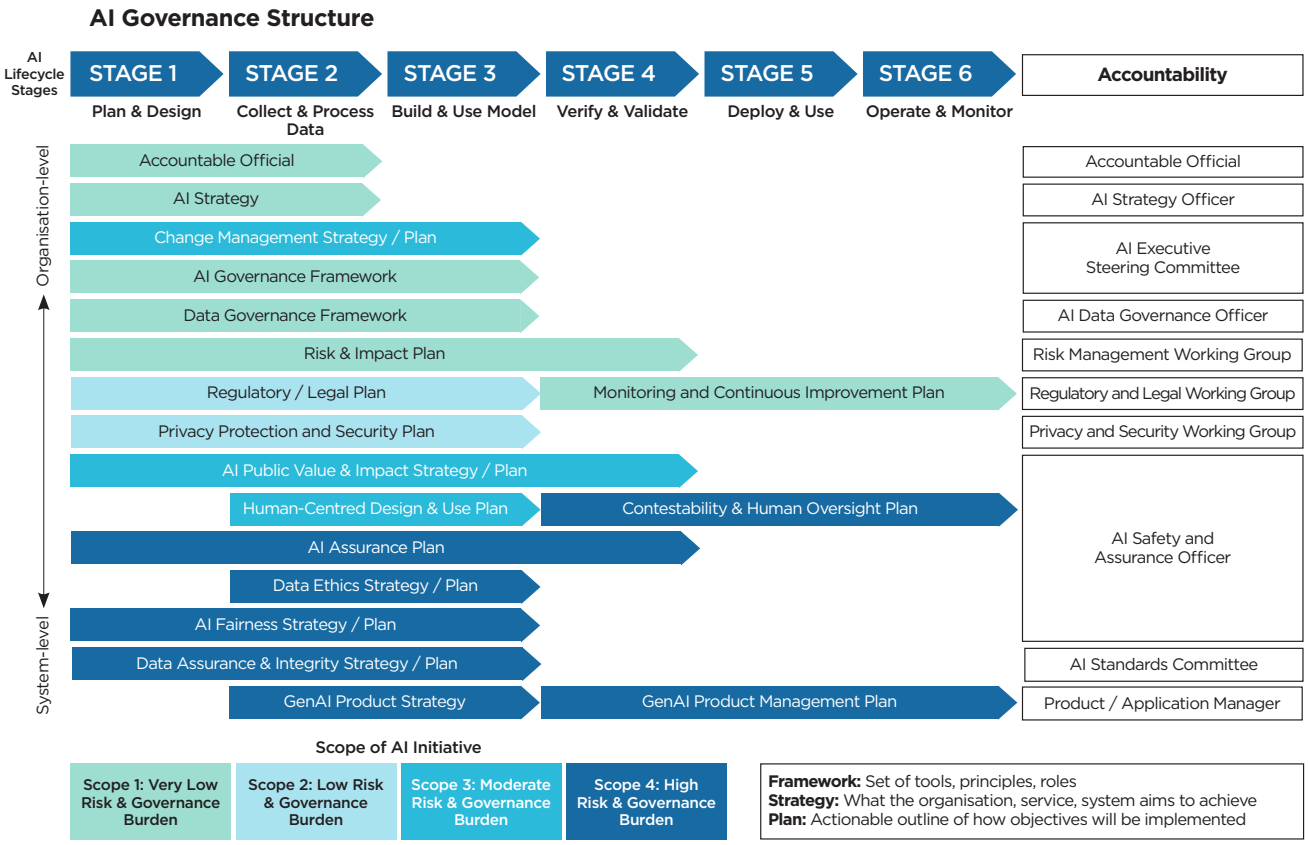

Figure 4: Representative AI governance structure for GSQ, mapping governance artefacts to scope tiers and lifecycle stages. Reproduced from Governance in the Age of AI: Readiness and Responsible Leadership (GSQ, 2025, Figure 5).

What good looks like

Leaders should expect the risk profile to shift across at least the following dimensions as an initiative moves from PoC towards production, drawing on the risk factors in the original white paper (GSQ, 2025, Section 4.3).

Capability and autonomy

- Type of AI service – from a low-risk consumer application (such as a simple chatbot) to a more complex self-trained model (such as an agentic system).
- Human-in-the-loop versus autonomous operation.

Data sensitivity

- AI access to public versus confidential, culturally sensitive or commercially sensitive data.
- Access control and disclosure obligations.

Connectivity and integration

- Isolation from sensitive or non-embargoed data systems.
- System dependencies and connections to backend or enterprise systems.
- Ability to invoke tools, APIs or workflows.

Human impact and reliance

- Internal versus external users, and how their security, privacy and trust are affected.
- Decision support versus decision-making.
- Regulatory consequences.

Supporting guidance

The Queensland Government FAIRA framework is one entry point for risk assessment in this context (Queensland Government, 2024). SCALE is deliberately framework-agnostic about the choice of risk methodology; the point is that risk should drive governance, not the other way around.

The concept of scaling governance effort according to the nature, scope and risk of the AI capability is consistent with industry approaches such as the AWS Generative AI Security Scoping Matrix (Saner and Lapidakis, 2023). While that matrix was developed from a technology and security perspective, the core principle holds: different AI systems carry different levels of risk, and as systems become more connected, autonomous, business-critical or externally facing, governance requirements should increase accordingly. Applying SCALE in this context ensures that high-level governance obligations can be easily understood and communicated, with a direct connection to the detailed guidance and frameworks that bring rigour to implementation.

Element	Primary (external)	Additional references (external)
A – Assess risk and align governance	Queensland Government – Foundational AI Risk Assessment (FAIRA); Australian Government – National Framework for the Assurance of AI in Government	Queensland Government – AI Governance Policy (2024); Privacy Act 1988 (Cth); Protective Security Policy Framework (PSPF); ISO/IEC 42005:2025; ISO/IEC 42001; ISO/IEC 38507; ACSC – Information Security Manual (ISM)

SCALE in practice: the GSQ Digital Librarian

Capability and autonomy

The Digital Librarian PoC was constrained to a low-autonomy role (Scope 1 or 2). While it used relatively sophisticated LLM and RAG techniques, it acted only as a retrieval and summarisation tool, leaving interpretation and decision-making to users. This limited autonomy kept governance proportionate during experimentation. Any future move towards greater automation (towards Scope 3 or 4) would require additional oversight.

Data sensitivity

The PoC operated on a curated subset of geological reports already publicly available through the GSQ Open Data Portal. This allowed GSQ to explore AI in a relatively low-risk data environment. Future expansion into embargoed, commercially sensitive or culturally sensitive information would materially change the governance and data stewardship requirements. Recognising this early, GSQ began considering the broader requirements associated with responsible AI adoption – including alignment with government AI ethics principles, relevant legislative obligations, and organisational data governance – as part of planning beyond the PoC.

Connectivity and integration

The PoC was deliberately isolated from enterprise systems and operated within a firewalled environment, specifically to reduce complexity and risk. A production service would require deeper integration with other Queensland Government systems, authentication services and business systems. Each additional integration would introduce new dependencies and governance obligations.

Human impact and reliance

While the Digital Librarian was designed for decision support, there is nonetheless the potential for users to place increasing trust in AI-generated responses as the capability matures. This reliance – whether perceived or actual – creates new expectations about transparency, explainability, accuracy and accountability, especially if outputs are used to support investment, regulatory or operational decisions. These changes illustrate a SCALE principle: governance should evolve as organisational reliance, connectivity and risk increase.

Leader's prompt

For your initiative right now: what scope tier is it operating at, and what scope tier are you funding it to reach? If those two answers do not match, governance and investment are misaligned.

L – Lift readiness and maturity

AI adoption quickly becomes a data, governance, readiness and leadership problem.

AI adoption in organisations quickly becomes a data, governance, readiness and leadership problem. Technical capability and organisational capability rarely mature at the same rate. This element of SCALE focuses on identifying the readiness gaps and understanding what must change before the next stage of maturity can be achieved.

What good looks like

Strategy

- Endorsed AI strategy with clear alignment to organisational priorities.
- A governance framework that scales with the AI lifecycle.
- Regulatory and legal obligations addressed.
- Ethical and human-centred values embedded.

Organisation

- Committed leadership sponsoring AI initiatives.
- Engaged stakeholders and active change management.
- A supportive, skilled and resourced workforce.
- Robust business and financial foundations.
- Cultural readiness and digital literacy.

Data

- High-quality, complete and reliable data with assurance processes.
- Provenance and integrity controls.
- Robust security and privacy safeguards.
- Processes to detect and mitigate bias, drift or modification.
- Lifecycle data governance.

Technology

- Agile, user-informed development with modular design and stage-gates.
- A clear product strategy.
- Robust AI model management.
- Fit-for-purpose, scalable, enterprise-aligned infrastructure.
- Service management plans for sustained performance.

Supporting guidance

The original white paper sets out a multi-domain AI Readiness Assessment Framework anchored in four domains: Strategy, Organisation, Data and Technology (GSQ, 2025, Section 3). These domains expand into twelve readiness areas, including Transformation Readiness, Cultural Readiness, Digital Readiness, Data Readiness, Strategic Momentum, Governance Capability, Digital Culture and Innovation Agility (GSQ, 2025, Section 3.1, Figure 2). Most organisations may find that one or two of these intersections explain their current sticking point more clearly than any single domain on its own.

The original white paper also presents a four-axis radar (spider) diagram that scores each domain from 0 to 100 (GSQ, 2025, Section 3.3, Figure 3). The diagram is most useful not as a score, but as a structured conversation. Three observations from the GSQ experience are worth highlighting for any leader using it:

- Strength can mask risk. A profile that is strong on Data and Technology but weak on Strategy and Organisation cannot be scaled on the strength of the leading domains alone.
- Profiles change with lifecycle stage. The target profile for an early-stage initiative is not the same as for one approaching production. Readiness is staged, not absolute, and is re-scored at each SCALE cycle.
- Low readiness with a high-risk use case raises governance. The readiness profile and the governance tier work together – low readiness combined with high exposure triggers higher governance requirements.

The original GSQ framework is complemented by a range of Australian and Queensland Government guidance that supports organisational readiness assessment for AI adoption, including the DTA Guidance for AI Proof-of-Concept to Scale, the AI Plan for the Australian Public Service 2025, the Australian Government Architecture, the Queensland Government Enterprise Architecture, and the Queensland FAIRA framework. These resources focus on organisational capability, workforce readiness, technology foundations and long-term operating maturity.

Element	Primary (external)	Additional references (external)
L – Lift readiness and maturity	Department of Finance, DTA & APSC – AI Plan for the Australian Public Service 2025; DTA – PoC-to-Scale Guidance, Appendix 1 (readiness)	Australian Government – Australian Government Architecture (AGA); Queensland Government – QGEA

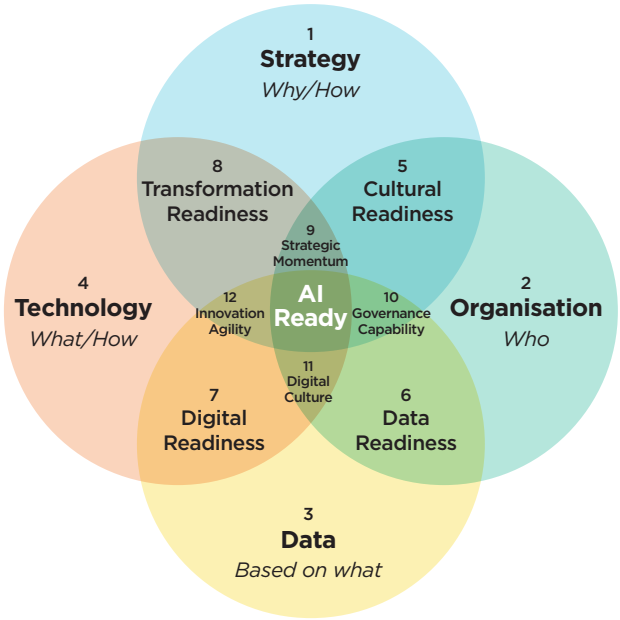

Figure 5: Multidimensional readiness framework offers leaders the opportunity to assess not only organisational readiness in the 4 key domains but also to assess readiness maturity across combinations of these domains. Reproduced from Governance in the Age of AI: Readiness and Responsible Leadership (GSQ, 2025, Figure 2).

SCALE in practice: the GSQ Digital Librarian

Strategy and Organisation

GSQ's initial readiness assessment showed that the Digital Librarian had benefited from years of prior investment in data and digital capability, but lagged on Strategy and Organisation. The PoC was delivered successfully within a small project team. However, scaling it raised questions about organisational ownership, workforce capability, change management, executive sponsorship and longer-term support models. Without uplift in the lagging domains, even a strong PoC cannot be scaled. These findings informed current work on a GSQ AI strategy and the establishment of approaches such as SCALE to support future decision-making.

Data

Data was one of GSQ's stronger readiness areas. Pre-existing investment in data management and digital platforms already provided access to data suitable for experimentation. The readiness assessment also highlighted an important distinction between being data-ready and being AI-data-ready. Scaling AI beyond a PoC would require additional controls around data provenance and the traceability and explainability of AI-generated outputs. This nuance is non-trivial and may require existing governance arrangements to evolve, or dedicated AI data governance responsibilities to be established.

Technology

The PoC demonstrated that contemporary AI technologies could deliver value quickly. It also highlighted that technology decisions suitable for rapid experimentation would need to be revisited before enterprise deployment, including enterprise architecture, system integration and government procurement pathways.

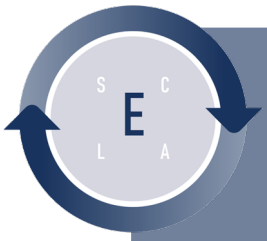

E – Embed, evaluate and assure

Every lifecycle stage ends with a decision, not a deployment.

AI services are not deployed once and left to run. They operate in environments where users, data, models, regulation, organisational reliance and risk all continue to change. Embedding, evaluating and assuring the service is what turns a successful deployment into a capability the organisation can stand behind over time.

The "how" of this element is less about a single event and more about an ongoing operating discipline. It has three connected practices: embedding the service into business operations so that it is owned, supported and used as intended; evaluating it continuously against the value it was meant to deliver and the risks it carries; and assuring that governance, controls and obligations remain fit for purpose as the service matures and reliance grows. Each lifecycle stage closes with an explicit decision – proceed, pause, redesign, scale or stop – informed by what evaluation and assurance have surfaced.

This is also the element that connects one SCALE cycle to the next. The outputs of Embed, evaluate and assure feed directly into the next cycle's Start (problem re-confirmation), Confirm (re-baselining), Assess (risk re-tiering) and Lift (readiness re-scoring).

Embed – operationalising the service

Embedding is the work of moving an AI capability out of project mode and into a managed operating state. It is the practical bridge between "we have built it" and "the organisation runs it".

- Service ownership and accountability assigned. A named business owner is accountable for the service, supported by the AI governance roles in the original white paper's accountability model (GSQ, 2025, Section 4.1 and Appendix B). The Accountable Official, AI Strategy Officer and AI Data Governance Officer roles are activated proportionate to the scope tier.
- Operating model in place. Funding, support, incident management, change management, user support and training pathways are defined and resourced – not aspirational. This aligns with ITIL service management practice (PeopleCert, n.d.) and with the GenAI Product Management Plan in the original white paper (GSQ, 2025, Section 4.2, Figure 5).
- Service integrated into the business workflow. Users have onboarded, expert-review pathways are in use, and the service has a defined place in the workflow it was designed to support rather than running adjacent to it.

- Records, transparency and disclosure obligations met. Responsible-disclosure notices are in place where users are interacting with AI, in line with the Policy for the Responsible Use of AI in Government and Queensland public records obligations.
- Retirement criteria defined. The DTA Technical Standard's Discover-Operate-Retire lifecycle makes retirement a first-class concern. The conditions under which the service would be paused, redesigned or retired are written down at the point it goes live, not invented later.

Evaluate – continuous measurement against intent

Evaluation is the discipline of checking whether the service is still delivering the value it was funded for, against the problem it was originally designed to solve. It runs continuously, with formal review at every lifecycle stage gate.

- Performance against the value hypothesis. The benefits articulated in Start (time saved, quality improved, risk reduced, decisions better supported) are measured and reported. If the value hypothesis is not being met, that is itself the evaluation finding.
- Model performance, drift and quality monitoring. Accuracy, answer quality, citation integrity, retrieval performance and drift are monitored on a defined cadence. This sits in the original white paper's Monitoring and Continuous Improvement Plan (GSQ, 2025, Appendix A) and aligns with the NIST AI RMF Measure and Manage functions and ISO/IEC 42001 performance-evaluation requirements.
- Data evaluation. Data quality, provenance, bias, completeness and access controls are reviewed against the criteria in the white paper's Data domain (GSQ, 2025, Appendix C.3) and the ACSC AI Data Security guidance. The distinction between data-ready and AI-data-ready, surfaced in Lift, is revisited here.
- User and stakeholder evaluation. User satisfaction, trust, behavioural change and any over-reliance signals are tracked – directly responsive to the risk-and-reliance prompt added in Confirm. Where contestability is in scope (Scope 3 and 4 services), user-raised challenges are logged and resolved through the Contestability and Human Oversight Plan (GSQ, 2025, Appendix A).
- Lessons captured and shared. Evaluation findings are recorded in a form that can inform the next SCALE cycle, not lost in operational noise. This is the mechanism by which a pilot cycle's evaluation becomes a production cycle's starting point.

Assure – independent confidence the service remains fit for purpose

Assurance provides independent confidence – to the Accountable Official, the executive sponsor, governance bodies and ultimately the public – that the AI service continues to operate within the risk tolerance, governance posture and policy obligations it was approved against. It is distinct from evaluation: evaluation asks "is it working?"; assurance asks "would an independent reviewer reach the same conclusion?"

- AI Assurance Plan activated. The AI Assurance Plan from the original white paper (GSQ, 2025, Appendix A) is implemented and reviewed at each lifecycle stage, positioned to align with the National Framework for the Assurance of AI in Government.
- Independent review of strategic and governance alignment. Assurance reviews check that the service still aligns with the AI Strategy, the Accountability Framework, regulatory and legal obligations, and the scope-tier governance artefacts identified in Assess. This is the function the AI Strategy Officer performs in the white paper's accountability chain (GSQ, 2025, Appendix B).
- Risk re-tiering when triggers occur. A defined set of triggers – material changes in users, data, model, integration, autonomy or external policy – initiates a re-tiering of the service's scope from 1–4 (GSQ, 2025, Section 4.3). Where a re-tiering moves the service up, the higher tier's governance artefacts must be in place before the change goes live.
- Alignment with recognised standards. Where the organisation has chosen to align, ISO/IEC 42001 (management system) and ISO/IEC 42005 (AI system impact assessment) provide an external benchmark for the maturity of assurance practice, with the NIST AI RMF Govern function as a complementary lens.
- Security, privacy and information assurance. Routine assurance against the ACSC Information Security Manual, the Protective Security Policy Framework, the Privacy Act 1988 (Cth) and the ACSC AI Data Security guidance is incorporated into the assurance cadence.

Decide – closing each lifecycle stage with an explicit decision

Each lifecycle stage closes with an explicit, recorded decision: proceed, pause, redesign, scale or stop. This is the practical mechanism that prevents a service from quietly drifting beyond the conditions it was approved against, and it is the entry point for the next SCALE cycle.

- The decision is explicit and recorded. The decision-maker (typically the Accountable Official or executive sponsor) records the decision, the evidence behind it, the conditions attached, and the next review point. This is consistent with Queensland's Project Assessment Framework and Business Case Development Framework gated investment expectations.

- The decision feeds the next SCALE cycle. A "proceed" or "scale" decision triggers the next SCALE cycle at the higher maturity stage (PoC → pilot, pilot → production). The five elements are revisited with the readiness and governance posture appropriate to that stage.
- "No decision" is treated as a finding. A service that is simply running without an explicit periodic decision is flagged as an assurance gap. This is the operational meaning of the leader's prompt: if it is simply running, evaluation has quietly stopped.

How the three practices reinforce each other

Embedding, evaluation and assurance are not sequential; they run in parallel and correct each other. Embedding without evaluation produces a service that is operated but not understood. Evaluation without assurance produces findings that are not independently confirmed and can be quietly de-prioritised. Assurance without embedding produces compliance artefacts that do not stick. The "how" of Embed, evaluate and assure is to keep all three in place at a level proportionate to the scope tier identified in Assess – and to close each lifecycle stage with the decision that the next SCALE cycle depends on.

PoC is technology-led. Post-PoC activity is strategy- and organisationally-led, and tied to the delivery lifecycle.

What changes after PoC

PoC stage	Post-PoC stages
Department of Finance, DTA & APSC – AI Plan for the Australian Public Service 2025; DTA – PoC-to-Scale Guidance, Appendix 1 (readiness)	Australian Government – Australian Government Architecture (AGA); Queensland Government – QGEA

What good looks like

Embed

- Assign clear service ownership and activate the AI governance roles proportionate to the scope tier.
- Stand up the operating model: funding, support, incident and change management, user onboarding and training.
- Integrate the service into the business workflow, with expert-review pathways in active use.
- Put responsible-disclosure, records and transparency obligations in place before go-live, and define retirement criteria up front.

Evaluate

- Review risk, readiness, maturity, data quality, user needs, governance obligations and operating requirements.
- Identify the gaps that matter most before the use case moves to the next lifecycle stage; focus first on the "must fix" issues.

Assure

- Provide independent confidence that the service still aligns with strategy, governance and policy obligations.
- Re-tier the service's scope when triggers occur, and confirm the higher-tier governance artefacts are in place.
- Run routine security, privacy and standards assurance on a defined cadence.

Decide and repeat

- Make an explicit stage-gate decision: proceed, pause, redesign, scale or stop.
- Feed lessons learned back into strategy, governance and future investment decisions.
- Begin the next SCALE cycle at the higher maturity stage, or retire the initiative.

Supporting guidance

Element	Primary (external)	Additional references (external)
E – Embed, evaluate and assure	DTA (2025) – Technical Standard for Government's Use of AI (Discover / Operate / Retire lifecycle)	Australian Government – Voluntary AI Safety Standard; NIST – AI Risk Management Framework; ISO/IEC 42001; ISO/IEC 5338; AXELOS – ITIL® (continual improvement)

SCALE in practice: the GSQ Digital Librarian

Evaluate

The first SCALE cycle confirmed that the Digital Librarian PoC could deliver value by improving access to geological knowledge. It also highlighted gaps around governance, data stewardship, ownership and organisational capability that require further evaluation before progressing to a pilot.

Embed

The Digital Librarian has not yet been embedded as a production service – and that is precisely where GSQ sits in the SCALE cycle. The PoC was deliberately operated as a bounded experiment rather than a managed operational capability, so the embedding work has been scoped rather than executed. GSQ has identified what embedding will require before the capability can be relied upon: a named business owner within GSQ's information services function; an operating and support model aligned to existing ICT service management; integration with the Open Data Portal and enterprise authentication; and responsible-disclosure and records arrangements consistent with Queensland Government obligations.

Recognising that embedding is an organisational rather than a technical step, GSQ is sequencing this work through its enterprise AI capability planning rather than rushing the Digital Librarian into operational use ahead of the readiness uplift identified in Lift. The first SCALE cycle therefore closes not with an embedded service, but with an explicit, evidence-based decision to prepare a pilot – the point at which embedding will begin in earnest.

Assure

Rather than moving directly from PoC to production, GSQ used the outcomes of the first SCALE cycle to prioritise organisational uplift. This work is currently in progress and includes development of a GSQ AI strategy, refinement of the AI readiness assessment methodology, and broader capability-building initiatives to support future AI adoption across the organisation.

Decide and repeat SCALE

The outcome of the first SCALE cycle was an informed decision to justify and prepare for a pilot. This is the point at which the SCALE lens loops back on itself. GSQ's Digital Librarian is now commencing a second SCALE cycle as part of pilot planning, with broader organisational change required to make that next step possible (GSQ, 2025, Section 2.2.2; GSQ Executive Forum, May 2026). Future transitions to production will repeat this process as the Digital Librarian capability, its users and organisational reliance continue to evolve.

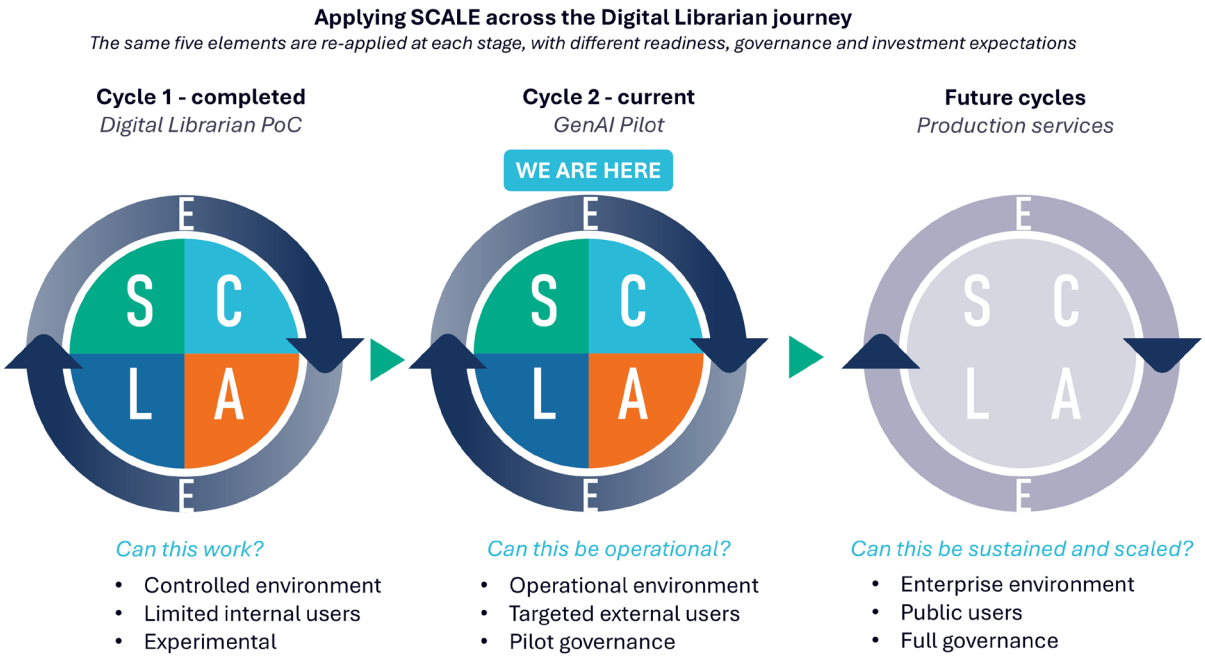

Figure 7: Applying SCALE across the Digital Librarian journey. The same five elements are re-applied at each stage, with different readiness, governance and investment expectations.

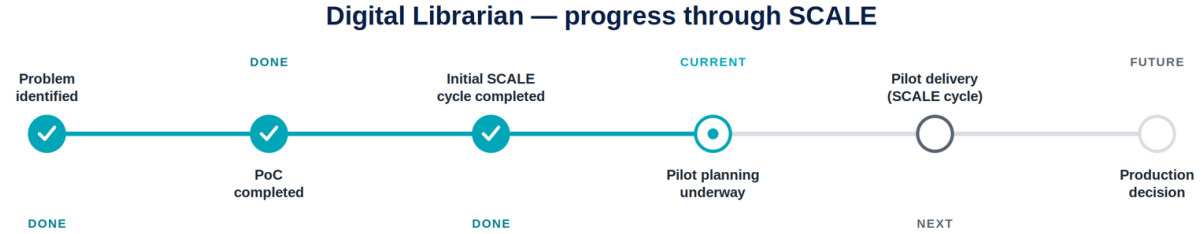

Figure 8: Digital Librarian – progress through SCALE.

References

Australian Cyber Security Centre (ACSC) (2025) AI data security. Australian Signals Directorate.

Australian Cyber Security Centre (ACSC) (2026) Information security manual. Australian Signals Directorate.

Australian Government (2026) Australian Government Architecture. Digital Transformation Agency.

Australian Government et al. (2024) National framework for the assurance of artificial intelligence in government, version 1.0, 21 June 2024. Australian Government.

Australian Government (2026) Protective Security Policy Framework. Department of Home Affairs.

Commonwealth of Australia (1988) Privacy Act 1988 (Cth). Federal Register of Legislation, latest compilation, accessed 14 July 2026.

Department of Industry, Science and Resources (2025) Voluntary AI Safety Standard. Australian Government.

Department of State Development, Infrastructure and Planning (2025) Business Case Development Framework. Queensland Government.

Digital Transformation Agency (DTA) (2025) AI Plan for the Australian Public Service 2025. Australian Government.

Digital Transformation Agency (DTA) (2024) Policy for the responsible use of AI in government, version 1.1. Australian Government.

Digital Transformation Agency (DTA) (2025) Technical standard for government's use of artificial intelligence. Australian Government.

Digital Transformation Agency (DTA) (2026) Guidance for AI proof of concept to scale. Australian Government.

Digital Transformation Agency (DTA) (2024) Digital Service Standard. Australian Government.

Gartner (2025) Lack of AI-ready data puts AI projects at risk, Q&A with Roxane Edjlali, 26 February 2025.

Geological Survey of Queensland (GSQ) (2025) Governance in the age of AI: Readiness and responsible leadership. GSQ and FrontierSI.

International Organization for Standardization (2022) ISO/IEC 38507:2022 Information technology — Governance of IT — Governance implications of the use of artificial intelligence by organizations.

International Organization for Standardization (2023) ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system.

International Organization for Standardization (2025) ISO/IEC 42005:2025 Information technology — Artificial intelligence (AI) — AI system impact assessment.

International Organization for Standardization (2023) ISO/IEC 5338:2023 Information technology — Artificial intelligence — AI system life cycle processes.

McKinsey & Company (2025) The state of AI in 2025: Agents, innovation, and transformation, 5 November 2025.

National Institute of Standards and Technology (NIST) (2023) Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1.

PeopleCert (n.d.) ITIL 4 Practitioner: Continual Improvement, accessed 14 July 2026.

PeopleCert International Limited (2023) Managing Successful Projects with PRINCE2® 7, 7th edn. PeopleCert International Limited.

Queensland Government (2024) Artificial intelligence governance policy, version 1.0.0, effective September 2024.

Queensland Government (2024) Digital investment principles, last updated 4 June 2024.

Queensland Government (2024) Foundational artificial intelligence risk assessment framework, version 1.0.0, September 2024.

Queensland Government (n.d.) Queensland Government Enterprise Architecture (QGEA), accessed 14 July 2026.

Queensland Treasury (2024) Project Assessment Framework: Policy overview, August 2024. Queensland Government.

Saner, M. and Lapidakis, M. (2023) 'Securing generative AI: An introduction to the Generative AI Security Scoping Matrix', AWS Security Blog, 19 October 2023.

Schuman, E. (2025) '88% of AI pilots fail to reach production—but that's not all on IT', CIO, 25 March 2025.

Wilkinson, M.D. et al. (2016) 'The FAIR Guiding Principles for scientific data management and stewardship', Scientific Data, 3, 160018.

Queensland Government

FRONTIER SI