

ASI VALUE DELIVERY SERIES 01

From AI opportunity to measurable value

A practical operating model for governed AI delivery

The ASI AI Value Delivery Framework

EXECUTIVE BRIEFING

AI value is a delivery discipline.

AI value does not come from the model alone. It comes from a disciplined operating system that connects diagnosis, economics, governance, engineering, adoption, and measurement.

Australia is still early in business adoption. In 2024-25, 12% of Australian businesses reported using AI. Only 7% measured the contribution of digital activities to overall business performance.² The figures measure different things, but together they expose the executive problem: access to technology is moving faster than the discipline required to prove value.

12%

Australian businesses reporting AI use in 2024-25²

7%

Businesses measuring digital contribution to performance²

6

Essential practices in current Australian adoption guidance¹

01**Start with an outcome**

Define the business decision, baseline, owner, and economic threshold before choosing a model.

02**Govern before release**

Design data boundaries, human control, evaluation, monitoring, and rollback into the system.

03**Measure in operation**

Compare live results with the agreed baseline, then improve, pause, or expand from evidence.

ABS adoption data does not measure maturity or return on investment. The measurement figure covers digital activities generally, not AI alone.

THE VALUE GAP

Access is not the operating model.

A model can produce a useful answer in minutes. A production workflow must produce the right output repeatedly, for the right user, from approved data, inside a controlled process, with an owner who can intervene. That is a different design problem.

The Productivity Commission identifies Australia's near-term opportunity primarily in adapting and implementing AI for specific local uses. Larger productivity gains depend on firms transforming core systems and adopting new tools, not merely gaining access to AI.³ ISO/IEC 42001 makes the same point from a governance perspective by treating AI as a management system that must be maintained and continually improved.¹⁰

Experiment	Production operating system
Can the model produce an answer?	Can the workflow improve a business outcome?
A convenient sample	Approved data with lineage, access, and retention
Anecdotal quality	Defined acceptance cases and regression tests
One enthusiastic operator	Named business, technical, and risk owners
A one-off demonstration	Monitoring, incident response, and rollback
Promising output	Measured value after adoption and operating cost

THE EXECUTIVE TEST

Can the organisation use the answer repeatedly, safely, and profitably?

If the answer depends on one person, one clean sample, or one untested assumption, it is still an experiment.


THE OPERATING MODEL

The ASI AI Value Delivery Framework

The framework uses six gates. Each gate must produce evidence, assign an owner, and end with a decision to proceed, pause, re-scope, or stop. This prevents urgency from silently removing delivery, security, privacy, or acceptance controls.

THE ASI AI VALUE DELIVERY FRAMEWORK

A gated path from opportunity to measured operation.



Every stage ends with evidence, an owner, and a decision to proceed, pause, or stop.

CONTROL

Govern across the lifecycle

NIST organises AI risk management around Govern, Map, Measure, and Manage.⁸

CONTEXT

Test the real workflow

Controlled tests do not reveal every risk that appears in operation. Monitor the live context.

CONTINUITY

Improve or retire

Responsible operation includes restriction, rollback, and decommissioning when evidence changes.^{1, 11}

STAGE 01 | VALUE

Make the investment case before the technology choice.

Start with one business outcome and one unit of work. Quantify the current volume, cycle time, labour, error, delay, revenue effect, customer effect, or prevented risk. Name the owner of every baseline and label every assumption.

THE ECONOMIC GATE

Value must be credible before the build is approved.

First-term commitment

1x

Minimum credible 12-month benefit

3x

Target credible 12-month benefit

5x

A defensible value case names:

- the baseline and its owner
- the measured business outcome
- the time window and assumptions
- the downside and operating cost

If the baseline cannot be validated, use operational acceptance metrics. Do not manufacture certainty.

ASI INVESTMENT DISCIPLINE

Minimum 3x. Target 5x.

Proceed when the credible 12-month benefit is at least three times the first-term commitment. Target five times. This is an ASI decision heuristic, not an industry standard.

Metric family	Examples
Business	Cost, revenue, hours, cycle time, error, experience, risk
Product	Adoption, completion, approval, time to output
AI	Evaluation score, correction rate, latency, cost per task
Operational	Uptime, failure rate, incidents, response time, monthly cost

STAGE 02 | DIAGNOSE

Map what actually happens.

The policy describes the intended process. The workflow reveals the delays, workarounds, handoffs, exceptions, and decisions that determine whether AI can create value.

OUTCOME

What must improve?

Define the business result, current baseline, target, time window, and acceptance owner.

WORKFLOW

Where does value leak?

Map triggers, actions, approvals, handoffs, exceptions, rework, delays, and final outputs.

SYSTEMS

What holds the truth?

Identify source systems, data owners, movement, freshness, quality, access, and retention.

DECISIONS

Where is judgement used?

Separate known rules from uncertain interpretation and define what requires human approval.

CONSTRAINTS

What must never happen?

Identify privacy, security, legal, procurement, safety, residency, and change constraints.

ADOPTION

Who must use it?

Name the operator, affected user, sponsor, technical owner, and escalation path.

READINESS GATE

Sponsor, access, and a bounded scope are mandatory.

A promising use case is not ready when there is no accountable sponsor, no representative data, no workflow owner, or no decision window. Pause and resolve the dependency before the delivery clock starts.

Current Australian Government guidance recommends documented accountability, business goals, acceptance criteria, impacts, risks, and review dates for each AI system.¹ For organisations subject to the Privacy Act, OAIC guidance also supports use-case-specific due diligence and privacy-by-design.⁵

STAGE 03 | PROVE

Use a bounded sandbox to remove uncertainty.

The proof is not a miniature production build. It is a controlled experiment designed to answer the highest-value uncertainty with the minimum necessary data, integrations, and operational risk.

01

Define the decision

State the hypothesis, input, expected output, business use, and what the proof must not do.

02

Approve the sample

Prefer a client-controlled environment. Minimise, de-identify, or redact data where practical.

03

Set acceptance cases

Include representative, difficult, unsafe, incomplete, and refusal cases before testing starts.

04

Measure the proof

Record quality, speed, correction, failure modes, latency, operating cost, and limitations.

05

Make the next decision

Recommend no build, a narrow workflow launch, further evidence, or a broader function design.

SANDBOX BOUNDARY**No production dependency. No irreversible action.**

A safe proof should not connect to production, create a live user dependency, or perform consequential actions without an approved control path.

Australian guidance recommends clear pre-deployment acceptance criteria that reflect intended use, context, and risk.¹ NIST's generative-AI profile adds governance, content provenance, pre-deployment testing, and incident disclosure as primary considerations.⁹

STAGE 04 | GOVERNANCE

Governance is part of the design.

Governance should accelerate a safe release by making ownership, constraints, evidence, and escalation explicit. It should not arrive as a final approval meeting after the workflow has already been built.

01

Accountability

Name the business, technical, data, privacy, security, and operational owners.

02

Impact planning

Document intended use, affected people, benefits, harms, limitations, and alternatives.

03

Risk management

Track cause, impact, likelihood, control, owner, residual risk, and approval.

04

Information sharing

Record system purpose, data sources, model and prompt versions, changes, and known limits.

05

Testing and monitoring

Define acceptance, regression, performance, security, and live monitoring evidence.

06

Human control

Design intervention, override, escalation, alternative paths, and decommissioning criteria.

AUSTRALIAN CONTEXT

Use current guidance as a practical baseline.

The National AI Centre describes these six essential practices as voluntary implementation guidance.¹ Commonwealth AI policy combines strategy, accountable ownership, registers, training, and impact assessment, but it is not a private-sector mandate.⁴

For APP entities, privacy obligations depend on the information flow and the use case. The OAIC recommends proportionate due diligence, privacy-by-design, careful testing, and responsible human verification for consequential AI-assisted decisions.⁵

STAGE 04 | ARCHITECTURE

The model is one component inside a production system.

PRODUCTION ARCHITECTURE

Define the operating system, not only the model.

01 User	Roles, interface, permissions, allowed actions
02 Application	Workflow, APIs, approvals, administration
03 AI	Tasks, models, tools, validation, prohibited behaviour
04 Data	Sources, stores, lineage, metadata, evaluation data
05 Integration	APIs, events, authentication, retries, failure handling
06 Security	Identity, least privilege, encryption, trust boundaries
07 Observability	Logs, quality, alerts, feedback, latency, cost
08 Deployment	Environments, release pipeline, backup, rollback

The model is one governed component inside a complete production system.

BUILD OR BUY

Keep standard work standard

Customise only where it creates the required outcome or a reusable operating advantage.

RULE OR MODEL

Separate certainty from judgement

Use deterministic logic for permissions, state, validation, and known calculations.

WORKFLOW OR AGENT

Earn the complexity

Use an agent only when constrained multi-step reasoning and tool use are required.

Australian cyber guidance treats security as a lifecycle requirement spanning design, development, deployment, and operation.⁶ Threat-model the full system, assess third-party dependencies, restrict allowed actions, and design monitoring for behavioural change.

STAGE 05 | BUILD

Build the simplest controlled path first.

Weekly increments reduce uncertainty only when each increment preserves the approved controls and produces evidence against the acceptance criteria.

01

Establish delivery controls

Repository, environments, ownership, reviewed changes, secrets management, and decision records.

02

Build one governed data path

Connect one source, preserve raw evidence where required, validate format, and handle failure.

03

Implement deterministic workflow logic

Define state, permissions, rules, approvals, retries, timeouts, and idempotency.

04

Add the AI capability

Constrain the task, model, tools, output schema, refusal behaviour, and validation.

05

Expose a usable operator experience

Show evidence, uncertainty, approvals, history, errors, and recovery paths.

06

Instrument before release

Track health, model quality, latency, cost, failures, user feedback, and business outcomes.

WEEKLY CONTROL LOOP

Plan. Build. Demonstrate. Record.

Demonstrate representative conditions, compare with acceptance criteria, update the risk register, and record the next decision. A demo is evidence, not theatre.


Do not wait until the end to add logging, human approvals, error handling, cost ceilings, or rollback. NIST recommends test, evaluation, verification, and validation across the AI lifecycle.⁸

STAGE 05 | RELEASE

Pass the seven-part production gate.

THE PRODUCTION GATE

No workflow goes live until all seven checks pass.



Release evidence must show what passed, who approved it, and what happens when conditions change.

A deadline is not a reason to hide residual risk. Exceptions require named acceptance and a recovery path.

Each pass needs evidence, an accountable approver, and a record of residual risk. Australian guidance recommends documented deployment authorisation, clear acceptance criteria, monitoring, and human control.¹ OAIC and ASD guidance reinforce privacy review, secure lifecycle controls, and careful testing before release.^{5, 6}

RELEASE RULE

A deadline does not remove a gate.

When an exception is unavoidable, document the residual risk, named acceptance, containment, monitoring, and recovery path.

STAGE 06 | OPERATE

Go-live starts the accountability clock.

Production introduces real users, changing inputs, third-party dependencies, new failure modes, and operating cost. Controls that passed in staging must be monitored in the context where the system now works.

Operating loop	Monitor	Decide
Health	Availability, integrations, queues, errors, response time	Restore, degrade safely, or roll back
Quality	Evaluation score, correction, refusal, drift, tool success	Tune, restrict, retrain, or change model
Cost	Usage, cost per task, provider spend, variance	Optimise, cap, route, or stop
Value	Adoption, throughput, cycle time, errors, business KPI	Improve, pause, or expand

CONTAIN

Make the system safe

Disable unsafe actions, isolate affected paths, preserve evidence, and protect users and data.

COMMUNICATE

Name the owner and next update

State impact, workaround, decision authority, and the next verified update time.

LEARN

Remove the repeat cause

Record root cause, corrective action, prevention, owner, and the test that proves the fix.

Australian cyber guidance recommends monitoring for sudden or gradual behavioural change.⁶ NIST and OECD likewise treat ongoing measurement, human oversight, override, repair, and decommissioning as lifecycle responsibilities.^{8, 11}

STAGE 06 | MEASURE

Expand from measured evidence, not momentum.

A production workflow earns the right to expand only after the organisation can show what changed, how the result was measured, what it cost, and which risks remain.

Evidence field	Required record
Baseline	Value, owner, date, data source, and confidence
Target	Business and operational acceptance criteria
Actual	Measured result over a defined time window
Adoption	Who used the workflow, how often, and where it was bypassed
Economics	Gross benefit, operating cost, change cost, and net value
Quality	Evaluation, correction, failures, limitations, and trend
Risk	Incidents, residual risks, control performance, and approval
Validation	Named business owner who confirms the evidence

01 No change
The workflow is stable, valuable, and needs no material intervention.

02 Optimise
Improve quality, cost, adoption, resilience, or operator experience within the current boundary.

03 Expand the function
Connect adjacent workflows through shared data, permissions, governance, and measurement.

04 Scale across functions
Reuse proven infrastructure and controls through separate, gated work packages.

60 TO 90 DAY REVIEW

Compare live results with the baseline.

Review value, adoption, incidents, quality, costs, risks, reusable assets, and remaining backlog. A recommendation to stop is a valid outcome.

ISO/IEC 42001 frames responsible AI as a management system that is maintained and continually improved.¹⁰ Australian adoption guidance also links business targets to system performance metrics and ongoing monitoring.¹

ILLUSTRATIVE EXAMPLE

Invoice exception handling

This is a hypothetical example. All volumes, costs, and results are invented to demonstrate the method. It is not an ASI client case study.

WORKFLOW
Accounts payable exceptions

Classify invoice exceptions, extract evidence, recommend a route, and require human approval before any system update.

BASELINE
8,000 invoices per month

Assume 15% need manual exception review, averaging 12 minutes at a loaded labour cost of \$65 per hour.

HYPOTHESIS
Reduce review effort

Assume 70% less review time on eligible exceptions and 80% operator adoption after stabilisation.

Step	Illustrative decision
Value	Annual review capacity is about \$187,000. Risk-adjusted benefit is about \$105,000 after the assumed reduction and adoption.
Economic gate	The capacity case alone supports a first-term commitment of no more than about \$35,000 at the ASI 3x threshold.
Diagnose	Map invoice sources, ERP state, approval rules, exception categories, duplicate controls, and audit requirements.
Prove	Use a de-identified set of 50 representative invoices, including difficult and unsafe cases, with no production connection.
Evaluate	Measure extraction accuracy, exception recall, evidence quality, correction, refusal, latency, cost, and failure modes.
Design	Keep ERP writes deterministic. Require human approval. Log the source, recommendation, decision, and final action.
Release	Proceed only after all seven production checks pass and the accountable finance owner accepts the residual risk.
Measure	Track adoption, review minutes, correction, duplicate rate, incidents, cost per invoice, and verified capacity released.

DECISION
The arithmetic can still say no.

If the verified baseline is lower, adoption is weak, integration is expensive, or risk controls absorb the benefit, the correct recommendation is to re-scope or stop.

EXECUTIVE READINESS SCORECARD

Is your organisation ready to prove one workflow?

Score each item 0, 1, or 2. Use evidence, not optimism. Sponsor, access, and governance cannot score zero if you intend to begin a bounded proof.

Criterion	0	1	2
Business outcome	Vague interest	Priority named	Metric and owner defined
Baseline	Unknown	Estimated	Verified and sourced
Executive sponsor	None	Interested manager	Accountable sponsor
Workflow boundary	Whole business	Function named	One workflow bounded
Data and access	Unavailable	Partial	Representative and approved
Governance	Unassigned	Policies identified	Owners and controls ready
Human control	Not considered	Manual review planned	Override and escalation designed
Acceptance	Demo standard	Some criteria	Test cases and thresholds approved
Operations	No owner	Support discussed	Monitoring and rollback owned
Value measurement	No plan	Target only	Baseline, target, cost, and review date

0 TO 8

Pause

Clarify the outcome, ownership, access, and risk boundary before spending on a proof.

9 TO 14

Diagnose

Run focused discovery to establish the workflow, baseline, data, controls, and decision.

15 TO 20

Prove one workflow

Proceed with a bounded sandbox only when the mandatory readiness conditions are present.

The scorecard adapts the accountability, impact, risk, testing, monitoring, and human-control themes in Australian guidance and NIST's Govern, Map, Measure, and Manage functions.^{1, 8}

THE NEXT DECISION

Prove the value before you scale.

A strong first engagement does not begin with a platform rollout. It begins with one function, one measurable outcome, one bounded proof, and a build-ready decision.

START WITH ONE FUNCTION

Book an AI Implementation Blueprint scoping session.

Map the workflow, quantify the case, identify the governance boundary, and prove the strongest opportunity in a controlled sandbox.

sameer@asiintelligence.ioasiintelligence.io

About ASI Intelligence

ASI Intelligence designs and delivers production-grade AI with measurable business value. The work is founder-led in Australia, with delivery partners across Asia-Pacific, and focused on mid-market financial services, insurance, manufacturing, and operationally complex businesses.

Sameer Chib

Co-founder, commercial and client outcomes

Ansh Mudgill

Co-founder, technical and engineering

Sources and notes

1. National AI Centre, Guidance for AI adoption: implementation guidance, Australian Government, 5 May 2026. ai.gov.au
2. Australian Bureau of Statistics, Characteristics of Australian Business, 2024-25, 25 June 2026. abs.gov.au
3. Productivity Commission, Making the most of the AI opportunity, 1 February 2024. pc.gov.au
4. Digital Transformation Agency, Policy for the responsible use of AI in government, version 2.0, 1 December 2025. digital.gov.au
5. Office of the Australian Information Commissioner, Guidance on privacy and the use of commercially available AI products, 21 October 2024. oaic.gov.au
6. Australian Signals Directorate et al., Guidelines for secure AI system development, 27 November 2023. cyber.gov.au
7. Office of the Australian Information Commissioner, Consultation on Guidance for Transparency in Automated Decision Making, 18 May 2026. oaic.gov.au
8. NIST, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023. doi.org
9. NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 26 July 2024. doi.org
10. ISO/IEC 42001:2023, Artificial intelligence management system, December 2023. iso.org
11. OECD, OECD AI Principles, updated 3 May 2024. oecd.org

This publication provides general information, not legal advice. Requirements depend on the organisation, sector, information flow, and use case. The ASI 3x minimum and 5x target are proprietary investment heuristics.