



Webinar

*by Dr Joe Sweeney*

A maturity-led framework for turning  
workplace AI into measurable value

# *Copilot: Where's the Value and How Do We Get It?*

# What We'll Cover Today

- Why Copilot value is hard to prove
- The hidden costs: verification tax and skill erosion
- A maturity-led roadmap for AI value
- New human-AI interaction models
- Measuring, assigning, and reporting benefits
- Governance, risk, and the 2026 regulatory horizon
- The technology group's role in AI innovation, benefits management, and literacy
- Your three-point action plan

# The Copilot Arrival Story

- **Bundled licences:**

Copilot 365 is frequently included in Microsoft enterprise agreement renewals

- **Entry fee:**

200 licences ≈ AU\$71,688 annually – and that is just the starting cost

- **Vendor claim:**

Microsoft promises ~1.2 hours saved per user per week at AU\$26.91 per user monthly

- **Real investment:**

Training, change management, governance, and data hygiene typically double the cost

- **Pressure drivers:**

APS AI Plan 2025, citizen service expectations, and staff demand

# The Human Landscape

The reality behind the enterprise agreement bundle

*Copilot:*  
Where's the  
Value and  
How Do  
We Get It?

# Why Value Is Hard to Justify

- **Diffused attribution:**  
Copilot is one input among many; benefits are hard to isolate
- **Non-financial returns:**  
Much of the value is quality, decision-making, and engagement – not hours
- **Workflow reality:**  
Initial drafting is only a fraction of total task time in regulated work
- **Verification burden:**  
Time saved in drafting is often consumed by checking AI output
- **Sunk-cost trap:**  
Without falsification criteria, organisations keep funding poor adoption

# The Verification Tax

- AI can reduce initial drafting time by 37–40%
- But drafting is only part of the work; the remaining ~50% is compliance, alignment, and verification
- **'Polished wrongness'**: AI output is grammatically perfect but can be subtly wrong
- The cost of finding and fixing errors is the verification tax
- Without workflow redesign, speed gains are lost to rework

# Polished Wrongness & Abdicated Co-Creation

## The Failure Mode

- Staff accept AI output without critical review
- ~40% of employees encounter plausible but incorrect AI outputs monthly
- Quality control switches off because the output looks 'done'
- This is abdicated co-creation

## The Antidote

- Train staff as editors-in-chief, not authors
- Build verification into workflows
- Reward error detection, not content volume
- Position AI as a capable but fallible intern

# The Skill Erosion Paradox

- Clinical studies show unaided diagnostic accuracy declines after just 3 months of AI reliance
- As work moves from 'doing' to 'choosing among AI outputs', domain expertise atrophies
- This creates sovereign risk: dependency on vendor black-box logic
- Organisations risk losing independent reasoning capability
- Consider 'AI-free' work periods and maintain human-verified decisions

# The Foundation for Value

Before you Scale, Fix the Foundations

*Copilot:*  
Where's the  
Value and  
How Do  
We Get It?

# Foundation First: Data Hygiene & Ground Truth

- **RAG is only as good as its sources:** Copilot's accuracy depends on the documents it can access
- **Oversharing risk:** Unmanaged SharePoint and OneDrive expose sensitive data to the wrong people.
- Duplicate and superseded documents reduce answer quality and increase hallucination
- **Recommendation:** Run a Purview workshop before broad deployment
- Limit access to well-governed, authoritative 'ground truth' repositories

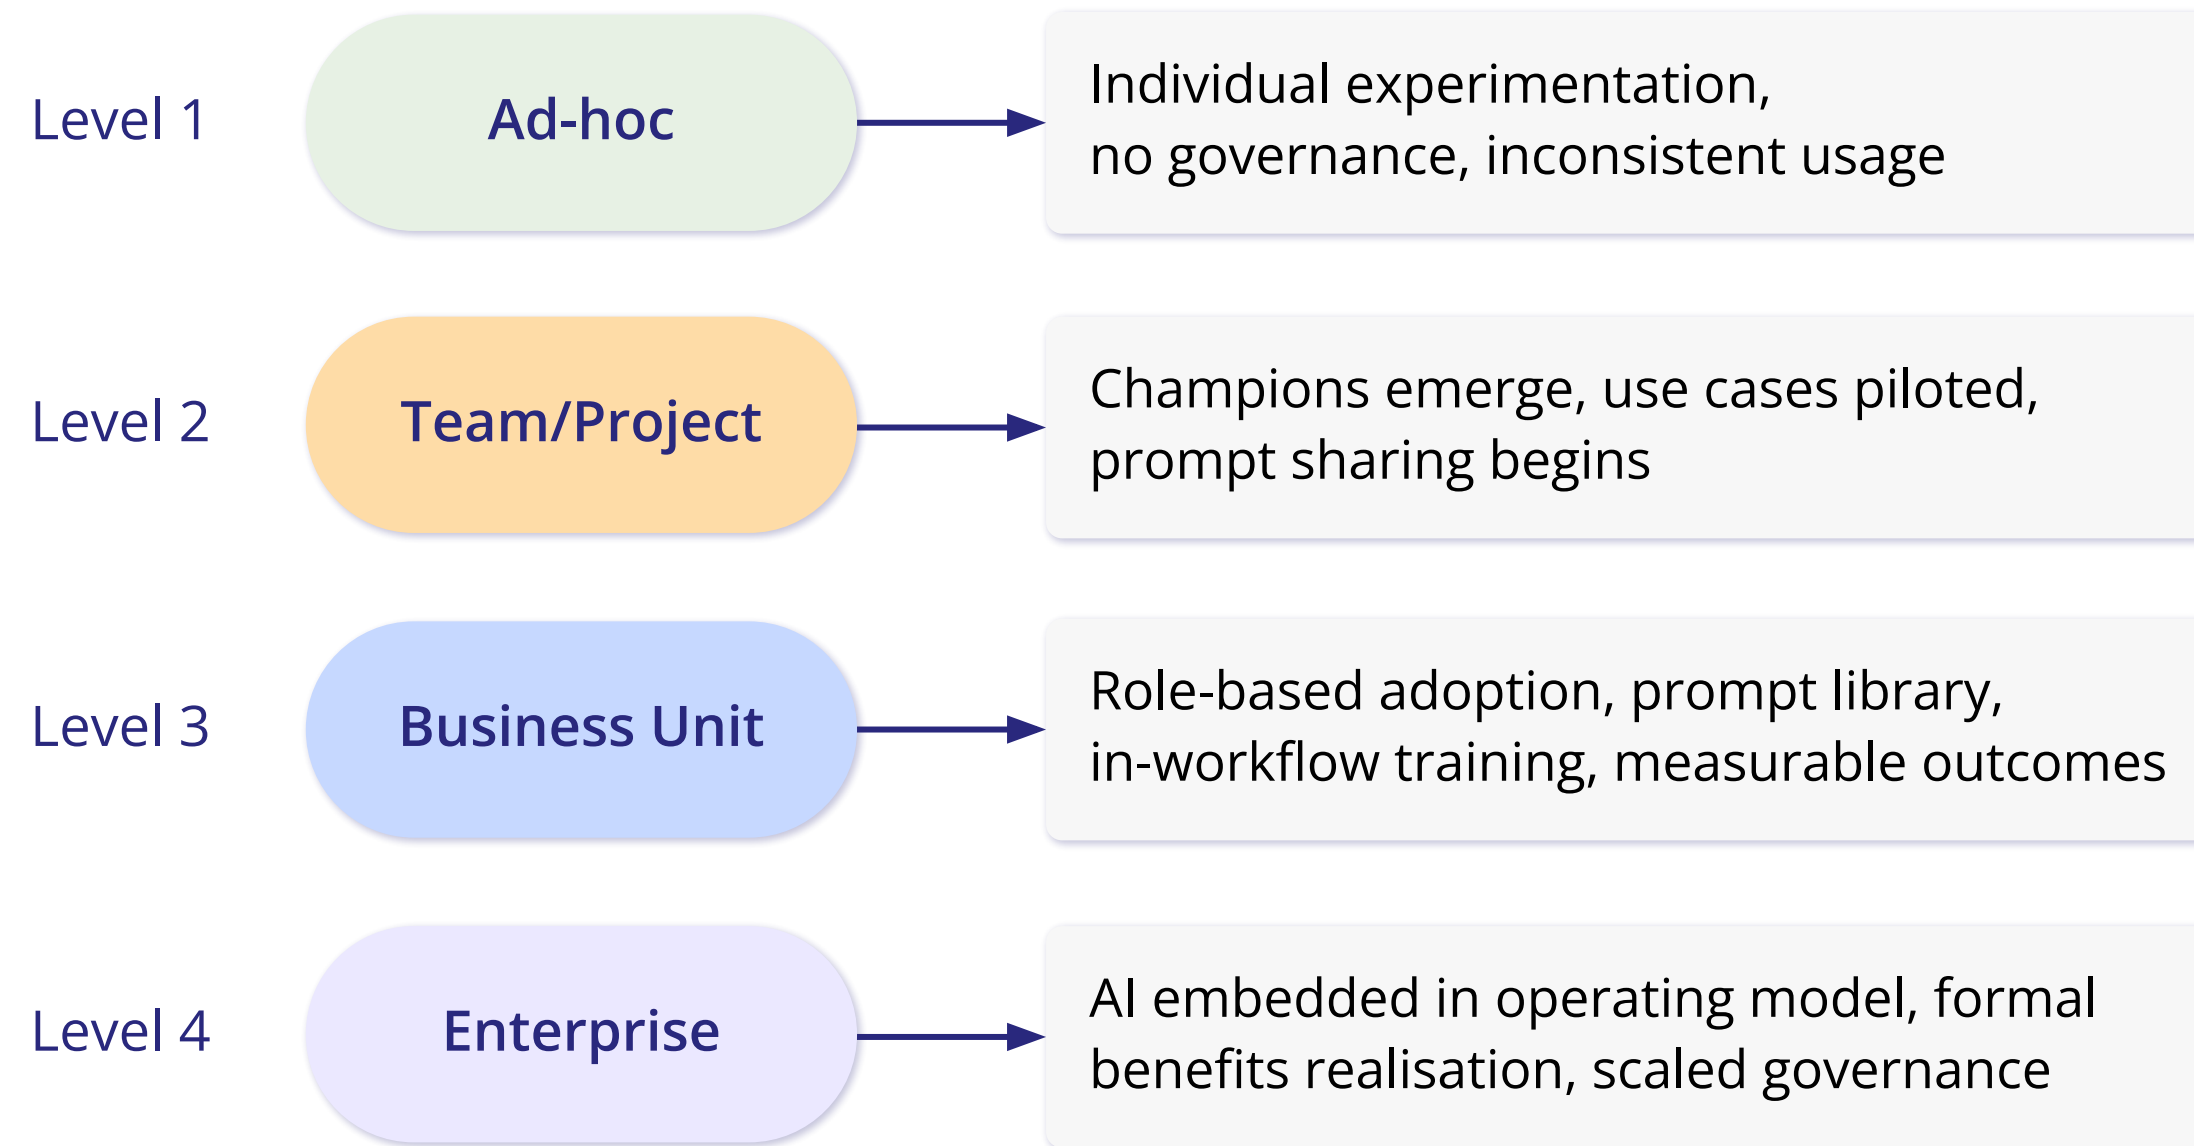
# The Falsification Rule

- Establish an irrefutable baseline before you start
- Define the criteria under which the project should be terminated or changed
- Example falsification criterion: <40% active usage of 3+ features weekly by day 90
- Telemetry options: Microsoft Viva Insights, focus hours, meeting loads, document revision cycles
- Protects against the sunk-cost fallacy common in digital transformations

# A Maturity-Led Roadmap

Move from Ad-hoc Experimentation to Enterprise AI Value

# The AI Maturity Model



*Financial benefits are only realistic once Level 3 is reached*

# *Phase 1: Foundation & Falsification (Months 1–3)*

- Quantify the 'before' state using telemetry and baseline audits
- Document the 10–15 highest-volume tasks in target business units
- Remediate oversharing and data quality issues
- Define falsification criteria and governance structure
- **Goal: An accountable, measurable program – not a technology rollout**

## *Phase 2: Scale & Adoption (Months 4–9)*

- Adoption is a training outcome, not a software switch
- Staff with 3+ forms of training report 75% higher confidence
- Move from generic 'lunch and 'earn' to role-based, in-workflow instruction
- Create a secure, curated Prompt Library on SharePoint
- Champions lead peer-to-peer workshops

# *Phase 3: Optimisation & Transformation* (Months 10–18)

- Redesign workflows to remove AI-revealed bottlenecks
- Conduct saved time audits to verify freed hours are reinvested
- Formalise benefit assignment in a benefits realisation register
- Use tangible performance data - not just surveys - to validate claims
- Examples: cases per employee, mean time to completion, error reduction, citizen satisfaction

# New Interaction Models

How Work Changes when AI Becomes a Teammate

# Centaur & Cyborgs

## The Centaur:

- Strategic division of labour
- Delegate reliable sub-tasks: summarisation, formatting, first drafts
- Human retains reasoning and proprietary judgment
- 'Directed co-creation'

## The Cyborg:

- Fused, continuous back-and-forth
- Human verifies and iterates in real time
- Best for tasks at the 'jagged frontier' of AI capability
- 'Fused co-creation'

# AI as a Capable but Fallible Intern

- Reframe Copilot from 'productivity tool' to 'highly capable but fallible intern'
- Sets realistic quality expectations
- Clarifies the manager relationship: the human is the editor-in-chief
- Applies to all workplace AI platforms, not just Copilot
- Helps overcome fear of deskilling and job displacement

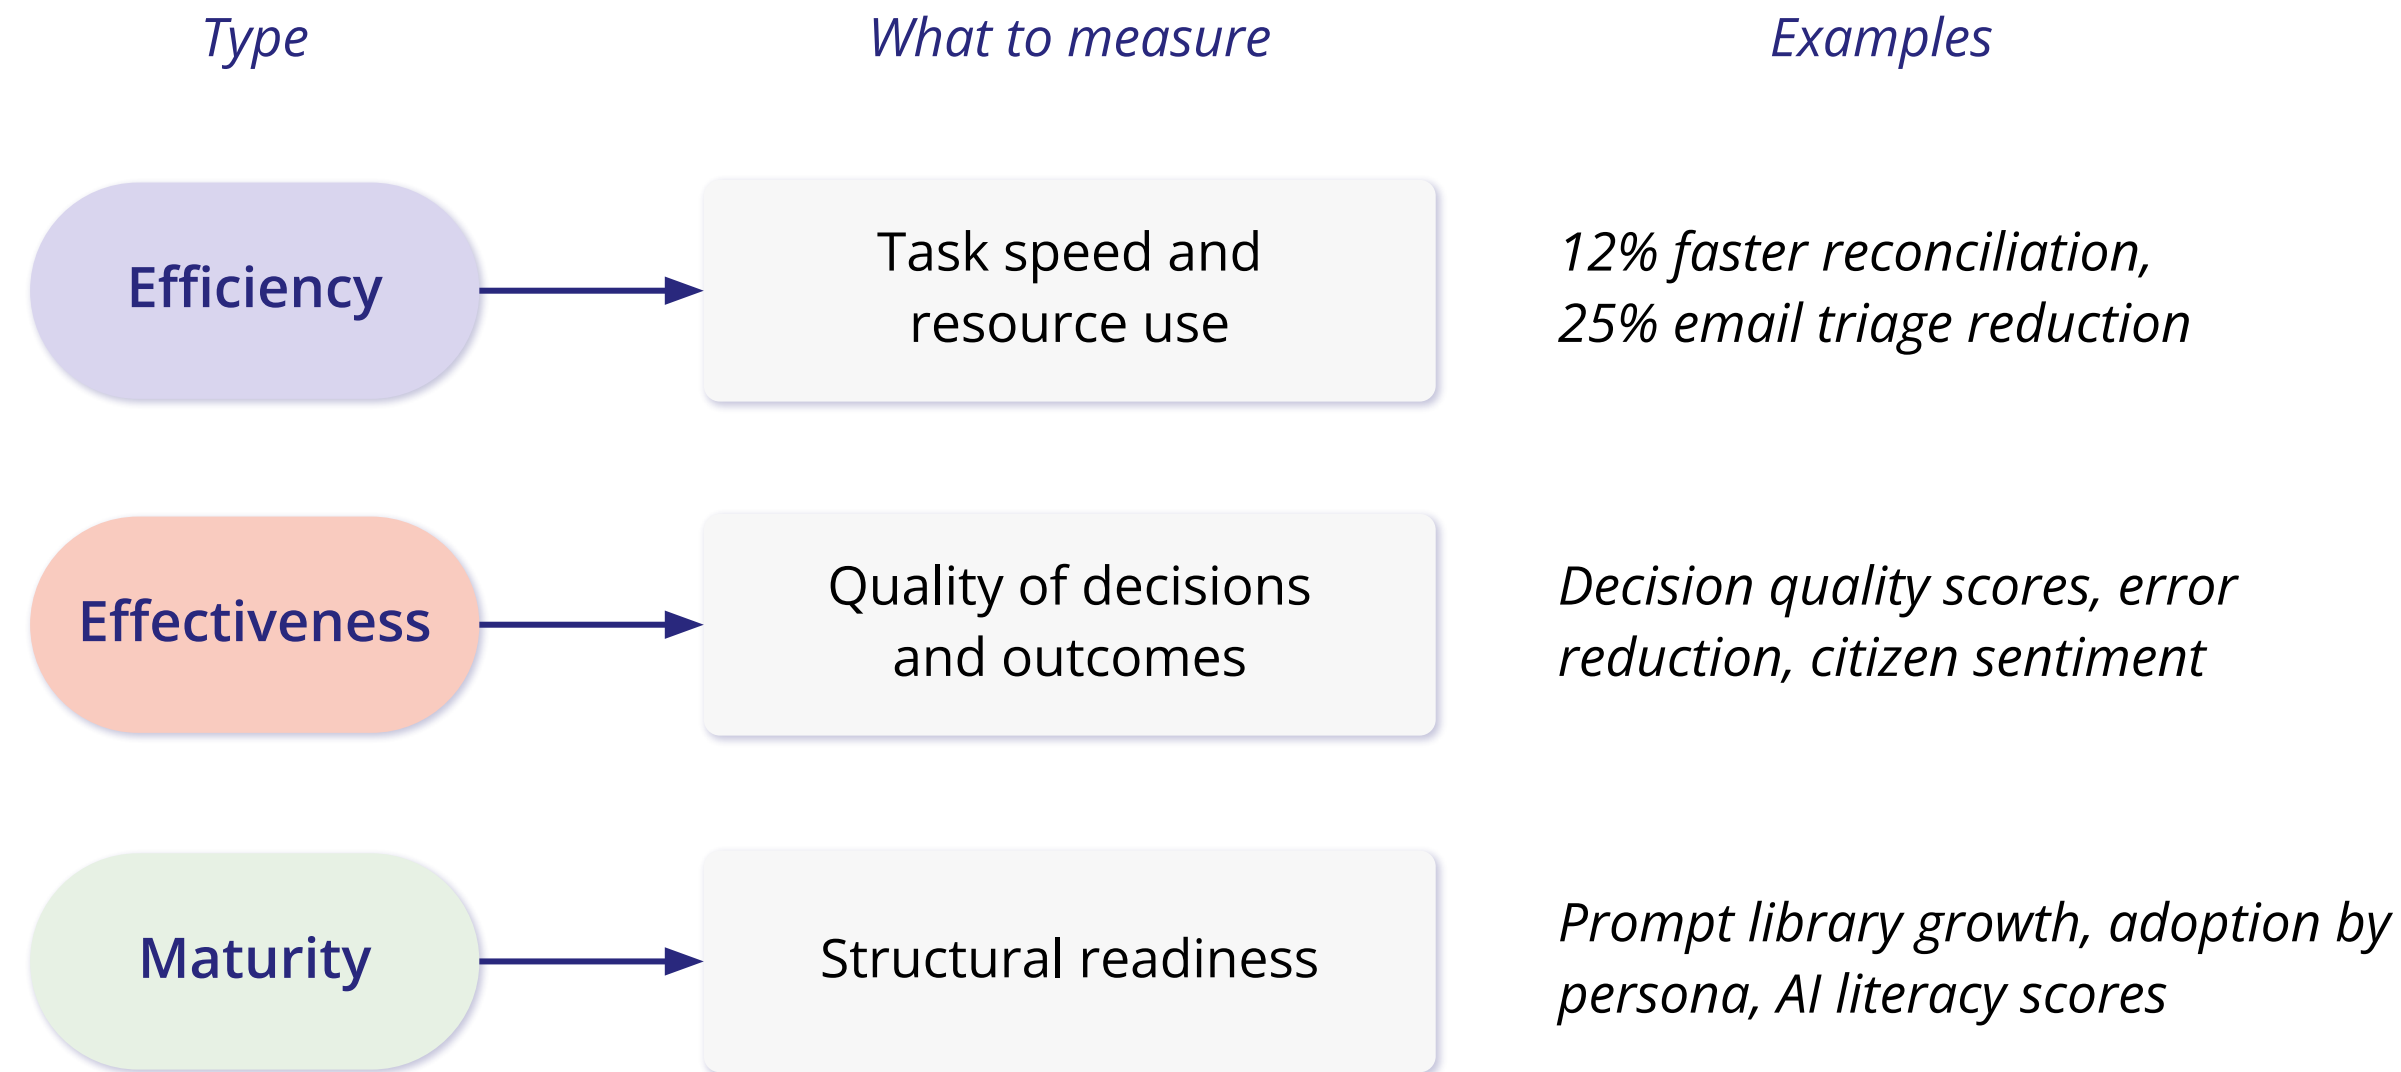
# Adoption Guilds & Champions

- Champions are early users who model effective usage for peers
- Guilds are cross-functional communities that identify lighthouse use cases
- Peer-to-peer learning outperforms top-down mandate
- **Examples:** City of Lahti and Oxfordshire County Council used champion networks to build momentum
- Lighthouse use cases must be driven by utility, not executive decree

# Measuring & Reporting Value

From Activity Dashboards to Benefits Realisation

# The Right Metrics: Activity vs. Outcomes



# The Ledger of Benefit: Total Cost of Ownership

<i>Cost Component</i>	<i>Annual Cost per User (AUD)</i>
<b>Copilot Licence</b>	\$432
<b>Implementation/Training</b>	\$216
<b>Governance &amp; Security</b>	\$144
<b>Total Fully-Loaded Cost</b>	\$792

# Risk-Adjusted NPV Scenarios (3-Year Horizon)

<i>Scenario</i>	<i>Probability</i>	<i>Year 1 NPV</i>	<i>Year 2 NPV</i>	<i>Year 3 NPV</i>	<i>Weighted Value</i>
<b>Pessimistic</b>	30%	-\$180	-\$120	-\$60	-\$108
<b>Base Case</b>	50%	+\$85	+\$220	+\$340	+\$322
<b>Optimistic</b>	20%	+\$420	+\$680	+\$890	+\$398
<b>Portfolio EV</b>					+\$612/user

*Copilot:*  
Where's the Value and How Do We Get It?

# Use Case Examples

- **Council Report Drafting:**  
4-hour task → 2.5 hours, saving 90 minutes per report
- **Planning Application Summarisation:**  
Extract constraints from developer submissions against local plans
- **Resident Enquiry Management:**  
Copilot Studio agents for bin collections, permits, 24/7 responsiveness
- **Meeting Administration:**  
Teams recaps, generating minutes and action lists
- **Contract Management:**  
Drafting templates and comparing vendor qualifications

# Governance & The Technology Group's Role

Why the Technology Group Must Lead —  
and Where it Must Step Back

# 2026 Regulatory Horizon

- **Privacy Act Reforms:** Transparency statements required for automated decision-making by December 2026
- **APS AI Plan 2025:** Mandates accountable officials and internal use case registers
- **Australian Consumer Law:** Misleading conduct and product liability for AI outputs
- **Contestability and Explainability:** Every AI-influenced decision must be traceable
- **Implication:** Governance must be designed in, not bolted on

# Accountability & RACI

<i>Task</i>	<i>Exec. Sponsor</i>	<i>CIO</i>	<i>AI Champions</i>	<i>CISO</i>	<i>Dept. Managers</i>
Strategy & ROI Planning	A	R	C	C	I
Data Hygiene & Labelling	I	A	C	R	I
Use Case Identification	I	C	R	I	A
Prompt Library Curation	I	I	R	C	A
Benefit Measurement	A	I	C	I	R
Governance Compliance	A	R	I	R	C

*Copilot:*  
Where's the Value and How Do We Get It?

# The Technology Group's Three Roles




## AI Innovation Broker

Identify, test, and scale  
use cases; connect  
business needs with AI  
capabilities



## Benefits Management Enabler

Build measurement  
frameworks, benefits  
realisation registers,  
and saved time audits



## AI Literacy Maturity Driver

Design training  
pathways, prompt  
libraries, and  
champion networks

# Your Action Plan

Three Practices to Start Immediately

# Three Definitive Practices

- **Forensic Measurement:** Add control-group experiments and saved time audits to activity dashboards
- **Structural Governance:** Align benefits realisation with the 2026 regulatory horizon; log and verify every AI-influenced decision
- **Maturity-First Investment:** Allocate at least 30% of the program budget to training, change management, and data hygiene

# Q & A



IBRS is a boutique Australian ICT Advisory Company.

We help our clients mitigate risk and validate their strategic decisions by providing independent and pragmatic advice while taking the time to understand their specific business issues.

© IBRS 2026

All Rights Reserved. This document and its entire contents may be used for information and educational purposes only.

All images remain the property of original copyright holders.