AI Governance:

What operational industries need to know.

How healthcare, manufacturing and mining organisations can adopt AI safely, securely and responsibly.

Brought to you by Colton Computers - your local technology partner in Orange, Bathurst and Dubbo.

A note before we dive in.

AI isn't coming. It's already here - and it's moving faster than most organisations have had time to prepare for.

If you're running a clinic in Dubbo, a processing plant outside Bathurst, or a mining operation in regional NSW, you're probably already feeling it. Staff are using AI tools. Vendors are bundling AI into their products. And somewhere in your organisation, someone is copy-pasting sensitive information into ChatGPT - without even realising it's a problem.

This guide isn't here to scare you. It's here to help you get ahead of it - practically, plainly, and without the tech jargon.



Section 1: Why AI is arriving faster than governance.

Australia's AI ecosystem is growing at a pace that's hard to keep up with. According to the Australian Trade and Investment Commission, more than 50% of Australian organisations report adopting AI in some form - and 49% of the general population is already using generative AI tools, a figure higher than both the US and UK.

The federal government is backing this hard. There's a $15 billion National Reconstruction Fund with $1 billion specifically allocated to critical technologies including AI, and a $392 million Industry Growth Program to support SMEs scaling up with new technology.

The momentum is real. But there's a catch - metropolitan organisations are adopting AI at a far faster pace than their regional counterparts. And in regional Australia, access to specialist legal and compliance expertise is often limited.

For operationally critical organisations, that gap isn't just a compliance headache. It's a real safety and security risk.



Section 2: The shared challenges facing regional healthcare, manufacturing and mining.

Different industries, same core problems.

Whether you're managing patient records, production lines, or extraction operations, regional organisations in healthcare, manufacturing and mining share a strikingly similar set of challenges when it comes to AI.

Operational continuity is non-negotiable. A hospital can't go offline. A manufacturing line that stops costs money by the minute. A mining operation has safety, logistical and financial consequences when systems fail. Unlike a city-based professional services firm, these organisations can't work offline for a day while they sort out a technology issue.

IT resources are limited. Regional businesses typically have smaller internal IT teams - or none at all. That means AI tools often get adopted without anyone evaluating the security implications. It happens fast, quietly, and with the best intentions.

Critical data is everywhere. Patient records. Manufacturing IP. Exploration data. Safety documentation. These are the kinds of assets that attract serious attention from cybercriminals - and they're increasingly being processed through AI tools without clear governance.

IT and Operational Technology (OT) are converging. This is a big one. In all three sectors, the technology that manages real-world processes - medical devices, supervisory control and data acquisition (SCADA) systems, fleet management software - is increasingly connected to broader IT networks. As the Australian Signals Directorate (ASD) notes, this convergence creates new attack surfaces that didn't exist a decade ago.

Compliance pressure is growing. Healthcare has privacy obligations around patient data. Manufacturing faces supply chain security expectations. Mining sits within critical infrastructure frameworks. Each sector faces its own regulatory landscape, but the underlying message is the same: you need to demonstrate that your technology risks are being managed.

Section 3: The cybersecurity risks of uncontrolled AI use.

Most AI-related security incidents don't come from sophisticated hackers. They come from well-meaning staff doing sensible things - just without the right guardrails in place.

Some of the most common risks we see include:

- **Sensitive data entered into public AI tools.** When a nurse uses a free AI assistant to help write clinical notes, or an engineer pastes process data into ChatGPT to troubleshoot an issue, that data may be used to train public AI models.
- **AI model drift.** AI systems can become less accurate over time as the environment they operate in changes. In a healthcare or OT context, degraded AI outputs can have real-world consequences - not just inconvenient ones.
- **Shadow AI.** Staff adopting AI tools independently, outside of IT oversight or approval, is increasingly common. You can't govern what you can't see.
- **Vendor-embedded AI.** Many OT vendors now bundle AI features directly into their products. These features may require internet connectivity, and their data handling practices aren't always transparent.
- **AI hallucinations in critical contexts.** AI tools can confidently produce incorrect information. In a safety-critical environment - think clinical decisions, equipment maintenance, or hazardous material handling - acting on a plausible-but-wrong AI output can cause harm.

The ASD’s guidance on AI in OT environments is clear: AI introduces risks that go beyond traditional cybersecurity, and organisations need specific governance frameworks to manage them safely.

Section 4: 10 questions every organisation should answer before deploying AI.

Before rolling out any AI tool - whether it's Microsoft 365 Copilot, a vendor-supplied analytics platform, or a standalone app - make sure you can answer these:

1. What data does this AI tool access, store, or transmit?
2. Where is that data stored - and is it kept in Australia?
3. Can staff use this tool with patient, customer, or operational data?
4. What happens to data entered into this tool? Is it used for model training?
5. Who in our organisation is allowed to use this tool - and for what purposes?
6. Does this tool have access to our operational systems, and if so, what level of access?
7. What happens if the AI produces an incorrect output - who is responsible?
8. Does our cyber insurance cover AI-related incidents?
9. Do our existing compliance obligations (privacy, safety, sector-specific regulation) apply to how we use this tool?
10. Have we updated our incident response plan to include AI-related failure scenarios?

If you're not sure of the answers, that's okay - that's exactly where a managed service provider (MSP) like Colton can help.

Section 5: Building an AI acceptable use policy.

An AI acceptable use policy doesn't need to be a 40-page legal document. It just needs to be clear, practical, and understood by the people it applies to.

A good AI acceptable use policy covers:

- What tools are approved for use, and the process for requesting approval of new ones
- What data can and cannot be entered into AI tools - particularly around patient, client, operational, and proprietary information
- Who is authorised to use AI tools and for which tasks
- How AI-generated outputs should be handled, including validation requirements before use in decision-making
- What to do if something goes wrong - reporting lines, escalation paths, and incident response steps

The Australian government's AI governance guidance recommends that organisations also define accountability clearly - someone needs to own this, and everyone needs to know who that is.

For smaller organisations without a dedicated IT or compliance team, an MSP can help develop, implement, and maintain this policy as part of a broader technology governance framework.



Section 6: Data classification and AI.

Not all data carries the same risk - but AI treats it all the same unless you set boundaries.

Data classification is the practice of categorising information based on its sensitivity, and then applying appropriate handling rules to each category. In an AI context, this means defining which types of data can be processed by which AI tools, under what conditions.

A simple three-tier model works well for most regional organisations:

Classification	Examples	AI use
Public	Marketing materials, general information	Approved for use with most AI tools
Internal	Internal communications, general business data	Approved AI tools with appropriate controls only
Confidential/Sensitive	Patient records, exploration data, production IP, financial data	Restricted - approved tools with strong data controls only (like Copilot with tenant-level data boundaries)

Copilot, for example, is designed so that it only accesses data within your organisation's Microsoft 365 environment. It doesn't send your data to external AI training datasets - your data stays within your tenant. That's a meaningful distinction from public AI tools, and it's exactly the kind of control that operationally critical organisations need.

The ASD specifically flags the risk of sensitive OT data being exposed through AI tools - including the risk of proprietary process data being embedded in an AI model where you may have no control over who accesses it or how it's used in the future.

Section 7: Human oversight and accountability.

AI is a tool. And like every tool, it needs a person responsible for how it's used.

The ASD is clear on this point: humans are ultimately responsible for functional safety. AI can assist, recommend, and automate - but in operationally critical environments, it should not make safety-critical decisions independently.

What does human oversight look like in practice?

- **In healthcare:** A clinician reviews and validates AI-generated documentation or diagnostic support outputs before acting on them
- **In manufacturing:** An engineer confirms AI-generated maintenance recommendations before scheduling downtime
- **In mining:** A site supervisor approves AI-assisted operational decisions before implementation

Beyond individual decisions, organisations should think about accountability at a structural level:

- Assign a named person (or team) responsible for AI governance
- Include AI tool usage in staff training and onboarding
- Build AI-related scenarios into existing incident response procedures
- Audit AI tool usage periodically to ensure policy compliance

One of the risks the ASD specifically calls out is **AI dependency and skill erosion** - the tendency for staff to stop developing manual skills because AI handles the task. In a sector where manual fallback capability might be genuinely safety-critical, this is worth taking seriously.

Section 8: AI governance checklist for regional organisations.

Use this checklist as a starting point for your AI governance review:

Foundation

- [ ] We have an approved list of AI tools that staff are permitted to use
- [ ] We have an AI acceptable use policy that has been communicated to all staff
- [ ] We have a named person or team responsible for AI governance

Data

- [ ] We have a data classification framework that defines what data can be used with AI tools
- [ ] We have confirmed how our approved AI tools handle, store, and protect organisational data
- [ ] Confidential and sensitive data is restricted from use with unapproved AI tools

Security

- [ ] We have reviewed our cybersecurity frameworks to include AI-specific risks
- [ ] Our incident response plan has been updated to include AI-related failure scenarios
- [ ] We have a process for evaluating and approving new AI tools before they are adopted

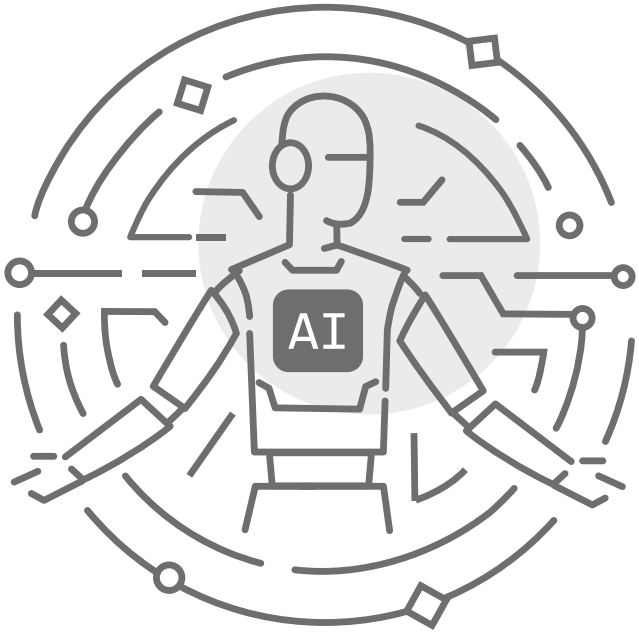
Section 8: AI governance checklist for regional organisations.

People

- [ ] Staff have received training on AI use within our organisation, including what is and isn't permitted
- [ ] We have clear escalation procedures for staff who are uncertain about AI use
- [ ] Human oversight is built into any process where AI outputs inform critical decisions

Compliance

- [ ] We have reviewed our sector-specific compliance obligations in the context of AI use
- [ ] We have confirmed whether our cyber insurance covers AI-related incidents



Section 9: How MSPs can support secure AI adoption.

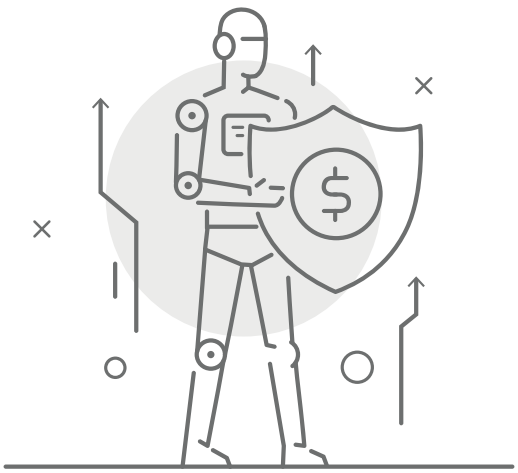
For most regional organisations, the reality is that you didn't get into healthcare, manufacturing, or mining to become an AI governance expert. You got into it to deliver great care, build great products, or run great operations.

That's where a local MSP comes in.

A good MSP - one that knows your business, your industry, and your region - can take on the heavy lifting of AI governance, so you can focus on what you do best. That means:

- **Evaluating AI tools** before they're adopted, so you're not flying blind on data security
- **Implementing enterprise-grade AI tools** with the right controls built in - like Copilot, configured correctly for your environment
- **Developing and maintaining AI policies** that are proportionate to your risk profile and compliance obligations
- **Training your team** on safe AI use, so the right habits are built in from the start
- **Monitoring and responding** to AI-related security events as part of a broader managed security service

Colton Computers has been helping businesses in Orange, Bathurst, Dubbo and the surrounding regions navigate exactly these kinds of technology challenges for years. We're not a faceless tech vendor - we're your neighbours, and we're invested in seeing regional businesses thrive safely.



Appendix: Ready-to-use AI guidelines template.

[Organisation Name] - AI acceptable use guidelines

Purpose:
To ensure that artificial intelligence tools are used safely, securely, and responsibly within our organisation.

Scope:
These guidelines apply to all staff, contractors, and volunteers who use AI tools in connection with their work for [Organisation Name].

Approved AI tools:
The following AI tools have been approved for use within our organisation: [List approved tools]

Permitted uses:
Approved AI tools may be used for: [List permitted use cases, like drafting internal documents, summarising non-sensitive reports, scheduling assistance]

Prohibited uses:
Staff must not enter the following types of information into any AI tool: [List restricted data types, like patient records, client financial data, proprietary process information, exploration data]

Data handling:
Before using any AI tool, staff must confirm that the tool has been approved and that the data they intend to use is permitted under this policy.

Validation of outputs:
AI-generated outputs must not be acted upon in a safety-critical context without human review and sign-off by an authorised person.

Reporting:
Any concerns about AI tool use, data exposure, or AI-related incidents should be reported to [Name/Role] immediately.

Review:
These guidelines will be reviewed [quarterly/annually] or when significant changes to our AI tool usage occur.

Want help putting this into practice? Colton Computers works with healthcare, manufacturing, and operational businesses across Central West NSW to implement secure, practical technology solutions - including Copilot with the right data controls for your environment.

Get in touch with the team at coltoncomputers.com.au or give us a call on 02 6361 1116. We'd love to have a chat.