

RLDatix

DatixCloudIQ (DCIQ)

Go beyond reporting risk and incidents with intelligent data and insights.

Learn More

✱ hello leaders

Saturday, August 22, 2026

NEWS60 SECOND SPOTLIGHTTOPICS ▼EDITIONSEVENTS

AI is different: Why your risk framework should be too

As AI adoption grows, treating it like traditional software leaves organisations vulnerable. Aged care providers must look beyond static risk registers to embrace continuous, dynamic governance for this fundamentally different class of technology.

Last updated on 25 July 2026

Image: Canva

The development and uptake of Artificial Intelligence (AI) systems continues to accelerate. In response to this business and technological sea change, regulators and governance bodies are now moving toward the idea that AI introduces not just an incremental improvement in existing tools, but a different class of organisational risk.

Meanwhile, most aged care providers are still trying to squeeze AI into existing software risk registers, which misses many of the important issues.

Ye olde risk register

It's a symptom of the modern world that processes that were state of the art five years ago will soon seem outdated, if not quaint.

Most organisations already have technology risk registers, software implementation procedures and cyber security plans. It's tempting to assume these are sufficient for AI systems – but they're not.

To see why, we need to understand the differences between traditional software and AI-based systems.

Why traditional software risk management worked

Traditional software introduced a structured system including:

- Analysts define requirements
- Developers write code
- Testers match performance to requirements
- Vendor releases a production version
- Users request changes
- Version updates are developed, tested and released.

Typical risks are usually:

- Bugs, including from untested edge cases
- Outages
- Security vulnerabilities
- Incorrect calculations
- Integration failures.

The critical point is that once a specific build is released, say Version 4.2, it remains Version 4.2 until someone changes it. What worked last week works today. "New" bugs were always there, just rarely triggered. Each change follows the same traditional software patterns from requirements, through to testing and release. Issues might arise, but the foundation doesn't morph beneath your feet.

Why AI is fundamentally different

Unlike conventional software, AI systems may produce different answers tomorrow even if the inputs are the same (or seem so to a human being). They can, without any explicit human involvement:

- Learn from changing information
- Access updated models with different "reasoning"
- Respond differently to identical questions, because of context
- Respond differently to similar questions, because of precise prompt wording
- Combine changing external knowledge
- Generate answers never explicitly programmed, via hidden "reasoning" that can't explain how conclusions were reached.

In effect, the software isn't changing but its *behaviour* is.

That distinction is incredibly important and has significant implications for those in your organisation who are responsible for risk management.

But it's equally important to note that these are not flaws in AI: not separate risks we have to eliminate. They are inherent to its function and value. It is *because* AI learns, adapts and correlates, *because* the models are frequently improved, that it can synthesise information and report linkages within minutes that a human would take hours or days to find, if they could at all.

Fundamentally, that's why AI presents a new class of risk: the source of its power is also the origin of its risks.

Derivative risks must also be considered

Despite common criticisms of AI, many people naturally trust its recommendations too readily: when the answers seem to appear instantly and almost magically out of a database too vast for a human to grasp, the temptation to just believe it is high.

But all users should realise that *no existing AI actually thinks*, reasons or understands anything at all. It merely generates outputs based on learned patterns, where those outputs are just more patterns, not representing conceptual knowledge within the AI itself.

On the other hand, anything new and powerful attracts regulators: aged care is in a highly regulated environment, and today's acceptable AI behaviour may not satisfy tomorrow's regulations.

The differences are summarised in this table:

Traditional Software	AI Systems
Static behaviour	Dynamic behaviour
Changes only when updated	Evolves over time
Testing before release	Continuous monitoring required
Defects corrected through code changes	Behaviour influenced by models, prompts and data
Predictable outputs	Probabilistic outputs
Version controlled	Model, prompt and data all influence behaviour
ICT governance	AI governance

What is AI Governance?

So if you need AI governance and it's fundamentally different, what is it?

Obviously you can't independently inspect and confirm all AI output: that would remove the benefits of using AI at all. Think of it like the taming of fire: the use of fire came with unique risks, but they had to be mitigated short of extinguishing the flames.

AI governance should include:

- Approved AI use cases
- Human oversight
- AI-specific risk assessments
- Monitoring outputs
- Validation of critical decisions
- Prompt management, including guardrails to minimise invention of poor guesses in the absence of actual evidence in your knowledge base.
- AI incident reporting
- Audit trails
- Periodic testing
- Supplier monitoring
- Staff education
- Version tracking.

Questions you should be asking

Instead of Information Governance asking "Is the software working?", consider AI Governance, where the questions you and your Board must ask are:

- Is the AI behaving as expected?
- Is it still trustworthy?
- Is it producing consistent outputs?
- Are people using it appropriately, including using approved AI tools rather than self-chosen consumer applications?
- Is it helping rather than harming?

Additional questions for your CIO and Board:

- Do we know every AI tool currently being used in our organisation?
- Have we assessed AI-specific risks rather than relying on our ICT risk register?
- Who is responsible for validating AI recommendations?
- How do we know when our AI behaviour changes?
- Could we explain an AI-assisted decision to a regulator or a resident's family?

It's becoming a mantra in the AI world that AI shouldn't replace people. Rather, it should free their time from rote tedium to enable what people are still uniquely good at: exercising judgement, empathy and accountability. This should become the basic principle supporting both your use of AI and how you mitigate its risks.

AI is moving fast – stay ready by adapting your systems

AI offers extraordinary opportunities to improve quality, efficiency and decision-making across aged care. However, organisations must recognise that AI can't be governed using yesterday's software management practices.

Effective AI governance requires continuous monitoring, appropriate human oversight and an understanding that AI systems evolve over time. Organisations that recognise this shift early and establish appropriate AI governance will be better placed to realise AI's benefits while maintaining safe, ethical and compliant care.

 **By Dr Robin Craig**
ThoughtWare/IonMy – Co-founder
"Dr Robin Craig is a scientist, philosopher and co-founder of ThoughtWare/IonMy, known for pioneering biotech research and challenging ideas."

Get the good stuff, weekly.

Get the news you need to your inbox.

Enter your email

GO

Sign up

Advertisement - See your product here!



ALTURA ESSENTIALS

Compliance training, reimagined

Streamline mandatory training with the 10-course compliance suite

Find out more →

AUG 15 – 21, 2026

• HOSPITALITY

Are we asking the wrong question about aged care kitchen shortages?

Lindsey O'Grady, Senior Catering Operations Specialist With over 75 per cent of aged care chefs reporting staffing shortages, benchmarked award rates may no longer be enough. Chef Lindsey O'Grady explores why providers must look beyond internal aged care pay comparisons to compete with the wider labour market.

• PALLIATIVE CARE

Palliative care isn't aged care's job, but the conversation is

Marion Piper, Contributor Victoria's first statewide report on palliative care reveals a stark gap between demand and access. As specialist shortages loom, Palliative Care Victoria CEO Violet Platt explores why aged care leaders must help bridge the gap through early conversations.

• INDUSTRY NEWS

Building the bridge to reshape First Nations aged care

Yvette Willison, Content Writer Interim First Nations Aged Care Commissioner Jodi Cassar brings more than two decades of public service experience to one of aged care's newest roles, sharing why listening, cultural safety and community partnerships are essential to creating meaningful, lasting reform for Aboriginal and Torres Strait Islander people.

✱ hello leaders

HELLO LEADERS MAG

Home

About Us

Contact Us

Make a Complaint

Advertise

Topics

Industry Updates

Events

Partners

Newsletter

Sitemap

© Hello Leaders Australia 2026