

AI Board Briefing: October 2026

October 2026 · Andrew France, CEO, SalesFix

Personal AI has been in the news this week, and it's turning up in new forms with the release of Meta's Muse and a slew of similar products, out now or coming soon. AI is arriving in organisations on people's bodies, not just their laptops. Display-less AI glasses with camera-and-audio functionality are going mainstream, and clip-on and pendant recorders capture meetings and calls and turn them into transcripts and notes.

The productivity upside is potentially great. However, staff and visitors will increasingly be recording, transcribing and capturing images in meetings and workplaces without any formal approval, and most organisations do not yet have policies that cover it.

Boards and their executive teams need to understand what is taking place and how it can affect their organisation.

The answer is to find out what's being used, give people a safe tool that's at least as good, and put the risk where the board and management can see it.

In this update we also cover the privacy policy changes you need to be considering, along with the AI risks that should be thought through.

Andrew France
CEO, SalesFix

In the News: Three Kinds of Personal AI

“Personal AI” now covers three different things. The first is a personal agent that acts on someone's behalf across their life. The second is staff using their own accounts with mainstream AI tools to do their work. The third is AI people wear: glasses, pendants and recorders that capture what happens around them. Each raises a different question for the board, and the lines between them are disappearing as the same tools take on more jobs.

	Personal AI agents	Personal accounts with AI tools	Wearable and ambient AI
Examples	Meta's Muse and the products that will follow it	A staff member's own Claude, ChatGPT, Gemini or Copilot account	AI glasses, pendant and clip-on recorders, AI notetakers in meetings
What it does	Remembers the person, connects to their email, calendar and apps, and takes actions for them	Answers questions, drafts, summarises and analyses when asked	Records, transcribes and captures images of whatever is around the wearer

	Personal AI agents	Personal accounts with AI tools	Wearable and ambient AI
How information gets in	The person connects a work inbox or calendar, or the agent remembers work conversations	The person pastes or uploads work material	It captures everyone in the room, whether they agreed or not
The main risk	An unapproved agent acting with access to your systems and data	Your data leaving your control, and unchecked output coming back	Recording clients, residents, staff and visitors without consent
The board question	What can it reach and do?	What goes in?	What does it capture, and about whom?

Personal Agents: The Line Between Work and Private Use Is Blurring

On 8 September Meta launched [Muse](#), which it describes as the first personal AI agent built for everyone. Muse doesn't wait for a question. It connects to a person's email, calendar and other apps, remembers details they've mentioned only once, and keeps working on their goals after they close the app. It can open a browser, fill out forms, negotiate and make purchases, checking with the person before sensitive actions like sending an email. It runs inside WhatsApp, is free for most uses, and is coming to Meta's AI glasses. Meta has since announced a version for small businesses.

Muse is rolling out in North America first. Users must opt out if they don't want their interactions used to train Meta's models, and Amazon has already [blocked it from shopping](#) on its site.

Muse isn't alone. OpenAI's Dot, Google's Gemini Spark, Grok Bot and at least four others are competing to be someone's personal agent. [Every's side-by-side comparison of eight of them](#) gives you further insight and concludes that the features have largely converged. They all connect to your apps and services, most can use a browser for you, and all offer customisable controls.

The more useful observation is about work and private life. None of these agents is strictly one or the other. They lean, with Muse pointing personal and Grok Bot and Dot pointing towards work, but each absorbs whatever context its owner gives it. A personal agent set up to juggle family bookings will just as happily read a work calendar, remember a client's name or draft a reply to a funder, if the person lets it. Today you can roughly tell an agent's leaning by whether it lives in WhatsApp and iMessage or in Slack and Teams.

The agents also differ on things a board should care about. Only two of the eight run on hardware the user chooses, and the rest run in the provider's cloud. One currently requires users to let it train on their data, while the others offer an opt-out. Some let users edit what the agent remembers, while others only let them delete everything.

Board Consideration: Personal Agents belong to the person, not the organisation, and it's useful precisely because it reaches across someone's whole life. The day an employee connects it to a work calendar or inbox, or lets it remember a client conversation, your data sits inside an agent you didn't approve and can't see. For most organisations the answer to "can staff connect a personal agent to a work account?" will be no for now. Your policies should be updated to reflect this.

Asking whether something is a work tool or a personal tool won't achieve much. The better question is what work context we've allowed these agents to absorb, and whether we could get it back out. That means writing your acceptable use policy around information: what may never be given to an agent the organisation doesn't control.

Personal Accounts: The Shadow AI Already in Your Organisation

This is the older problem, and it's bigger than most boards realise. Employment Hero's AI Paradox at Work report, a survey of 1,634 Australian workers and 1,008 business leaders, found one in three workers use AI without their employer's knowledge. A separate PagerDuty survey found 70% of Australian office professionals had used AI at work while believing it was against company policy, and 83% believed leadership plays by different AI rules ([HRD Australia, July 2026](#)).

Salesforce's own research points the same way. In a YouGov survey of 1,293 Australians commissioned by Salesforce, 56% of workers who use AI on the job said they rely on tools their employer hasn't provided or approved ([SecurityBrief Australia, June 2026](#)). In the interest of disclosure, SalesFix is a Salesforce partner.

On 7 September the UK's National Cyber Security Centre published [The hidden risks of shadow AI](#). Its advice is to reduce the risk rather than assume it can be eliminated, by understanding why staff reach for unapproved tools and giving them secure alternatives.

Board Consideration: "we've blocked ChatGPT" is no longer a satisfactory answer from management. The better answer is "here's what people use, here's the approved alternative, and here's how we know it's working". And the 83% figure should give directors pause: if your people think the rules are different for leaders, ask how board members are using AI with board papers.

Before buying anything new, check what you already license. Salesforce made Slackbot generally available in January as a personal agent for work, built into Slack and covered by Slack's security and compliance standards. Microsoft 365 Copilot and Gemini in Google Workspace offer similar capabilities inside the platforms many organisations already pay for. A personal agent the organisation controls is the most realistic answer to the ones it doesn't.

Wearables and Recorders: AI That Captures Everyone in the Room

On 28 July Kmart released its \$89 Anko Smart Glasses, with a built-in camera that records high-definition video and photos. They sold out across Australia within days. Meta's AI glasses keep getting more capable, pendant and clip-on recorders turn every meeting or call into a transcript, and AI notetakers now join video calls on behalf of staff and the people on the other side.

On 17 September the federal government said it was considering banning camera-equipped smart glasses in Commonwealth workplaces, which could be the first ban of its kind ([Reuters via The Edge](#)). The Attorney-General has asked the OAIC to investigate the privacy risks of smart glasses, including harassment and surveillance ([SmartCompany](#)), and some workplace lawyers are already advising employers to ban them outright ([Lawyers Weekly](#)). There's no single national law on recording devices. Each state and territory runs its own surveillance devices regime, and the new statutory tort for serious invasions of privacy adds another layer.

Board Consideration: the first two kinds of personal AI put your data at risk through your own people. This one captures information about people who never agreed to anything: clients, residents, colleagues and visitors. That makes it a consent and duty-of-care question as much as a data one. Your policy needs to say who may record, where, and with what notice, and it needs to cover visitors and contractors as well as staff. If you run aged care, disability or health services, settle your position on recording in care settings now, before a family member arrives wearing glasses.

More forms are coming, including AI browsers that act inside logged-in accounts and AI built into phone operating systems. The same three questions still apply: what goes in, what it can reach, and what it captures.

Risk Watch: Personal AI

Personal AI touches more board committees than most directors expect. Each risk below is tagged with the group that should own it. If you don't have that committee, the board owns it. The last column is for the executive team: if the board asks the question, this is what you want ready.

Risk	Owner	The board asks	Management prepares
Personal information about clients, donors, residents or staff entered into personal AI accounts	Audit and risk	Do we know what personal information has gone into tools we haven't assessed?	A short staff survey (anonymous, no consequences) and a list of the tools in use
Confidential or contracted data shared in breach of client, funder or partner agreements	Audit and risk	Do any of our contracts restrict sharing data with third-party tools?	A check of the top ten contracts for data-handling clauses
AI agents connected to work email, files or systems with no oversight	Audit and risk	What can an AI agent reach in our systems today, and who approved it?	A list of third-party apps connected to Microsoft 365, Google Workspace and the CRM
AI shaping decisions about people (eligibility, triage, hiring, credit, pricing)	Whole board	Is any decision about a person being made or shaped by software?	A first pass at the decision map needed for the 10 December privacy obligation
Staff hiding AI use, or believing different rules apply to leaders	People and culture	Do our people know what's allowed, and do leaders follow the same rules?	A one-page approved-use guide and an approved tool that's at least as good
Staff or visitors recording meetings, client sessions or care settings with AI glasses, pendants or notetakers	People and culture, with clinical governance in aged care and health	What's our position on recording in our workplaces and care settings, and does it cover visitors?	A recording policy covering staff, visitors and contractors, with notice and signage where needed
Care notes, assessments or clinical records drafted with unapproved AI	Clinical governance (aged care and health)	Where is AI touching care documentation, and who checks it?	Spot audit of a sample of care notes and assessments
Directors putting board papers into personal AI tools	Chair and company secretary	How are we, as directors, using AI with confidential board material?	A board protocol: which tool, which papers, and what never goes in

Worth Reading: AI Blunders That Caught the Regulators' Eye

None of these organisations set out to misuse AI. In every case a person trusted an output they didn't check, or used a tool in a setting it wasn't approved for. That's exactly the pattern shadow AI creates, and Australia already has a track record of regulators, courts and parliament responding to it.

Case	What happened	Who responded	The lesson for your board
Victorian child protection report (2024)	A child protection worker used ChatGPT to draft a report for the Children's Court. The report misrepresented a serious risk to a child as a positive, and the worker may have used it on around 100 other documents	The Victorian Information Commissioner investigated and issued a compliance notice requiring the department to ban generative AI for child protection staff	When frontline staff use personal AI on sensitive client records, the regulator's answer can be a blanket ban. For-purpose organisations working with vulnerable people are most exposed
Deloitte's report for DEWR (2025)	A A\$440,000 government review contained fabricated academic references and a made-up quote from a Federal Court judgment. Generative AI use was disclosed only in the corrected version	The department secured a refund of more than A\$97,000, and senators questioned who was actually doing the work being paid for	Ask your advisers and suppliers how they use AI on your work, and make disclosure a contract term. If it can happen to a Big Four firm, it can happen in your finance team
Lawyers filing fake case citations (2024 to 2026)	Lawyers in Victoria, NSW, WA, Queensland and SA filed submissions containing cases that don't exist. Australian courts have now seen more than 20 such incidents	Courts referred practitioners to legal regulators, and a Victorian lawyer became the first in Australia to be formally sanctioned	Professional duty can't be handed to a tool. The same applies to your clinicians, accountants, case workers and anyone else whose sign-off others rely on
Bunnings facial recognition (2024 to 2026)	Bunnings used AI facial recognition in stores to protect staff from violent customers	The OAIC found it breached privacy law. In February 2026 the Administrative Review Tribunal affirmed the findings that Bunnings breached transparency and notice obligations	A legitimate purpose doesn't excuse a lack of transparency. This is the same principle behind the 10 December automated decision-making obligation

The common thread is that no new AI law was needed. Privacy law, professional duties, consumer law and plain contract terms did the work. Your organisation is already exposed to every one of them.

Regulatory Countdown

The date that matters most is 10 December 2026, when the automated decision-making transparency obligation starts. It's just over two months away.

Date	What happens	Who it affects
30 November 2026	The federal Joint Select Committee on Artificial Intelligence is due to report	Everyone: a signal of where AI law is heading
10 December 2026	Privacy policies must disclose where a computer program makes, or substantially and directly shapes, decisions that significantly affect people. It covers assisted decisions, not only fully automated ones, and applies to decisions made from that date even if the system is older	Organisations covered by the Privacy Act. Most with turnover above \$3 million are, and health service providers are covered regardless of size
10 December 2026	The OAIC's Children's Online Privacy Code is due to be finalised. The OAIC has said it will apply to AI-enabled services such as chatbots	Organisations whose online services children are likely to use

Date	What happens	Who it affects
Now	ASIC told banks on 30 September that it will formally review their AI use in dealing with customers	Financial services, and anyone who expects their regulator to follow ASIC's lead
Early 2027	Legislation for the Australian Standards for AI, announced in July alongside a new Office of AI in Prime Minister and Cabinet, is expected in Parliament	Everyone, details to come

If you do one thing before Christmas, ask management for a list of the decisions about people that software is already shaping. You'll need it for 10 December, and it's also the start of a proper AI register.

Thirteen AI Risks for Your Risk Register

Most risk registers we see have one line for AI, usually under cyber. That's not enough to govern it. Here are thirteen entries you can copy into your register, rate against your own likelihood and consequence scale, and assign. Not every one will apply to you, and that's a useful conversation in itself.

#	Risk	What could go wrong	Controls to consider	Suggested owner
1	Unapproved AI use	Staff put personal, confidential or commercial information into personal AI accounts outside our control	Anonymous usage survey, approved alternative tool, acceptable use guide, browser and app monitoring	Audit and risk
2	Privacy and automated decisions	We breach the Privacy Act, including the transparency obligation for automated decisions from 10 December 2026	Map decisions about people that software shapes, update the privacy policy, privacy impact assessments for new AI	Audit and risk
3	Relying on wrong AI output	Inaccurate AI output ends up in board papers, advice, funding applications, care notes or client deliverables	Human review rules by document type, disclosure of AI use, spot audits	Executive, with clinical governance where relevant
4	AI agents and connected apps	An AI agent, including a staff member's personal agent, or a connected app with broad access is misused or compromised, exposing data or taking unauthorised actions	Register of connected apps, least-privilege access, approval before agents connect to core systems	Audit and risk
5	AI-enabled fraud and impersonation	Deepfake voice, video or email impersonating executives or suppliers leads to a fraudulent payment or data release	Call-back verification for payment and bank detail changes, staff awareness, finance controls	Finance and audit
6	AI inside existing software	Vendors switch on AI features in systems we already use, processing our data in ways we haven't assessed	Review vendor AI terms at renewal, default settings check, data processing agreements	Executive
7	Contract and confidentiality breach	Using AI on client, funder or partner data breaches contract terms on confidentiality or data handling	Contract clause review, data classification, approved-tool rules by data type	Executive

#	Risk	What could go wrong	Controls to consider	Suggested owner
8	Unfair or biased outcomes	AI-shaped decisions about eligibility, hiring, pricing or service disadvantage some groups, creating discrimination and reputational risk	Testing before go-live, human review of adverse decisions, complaints monitoring	Whole board
9	Workforce and culture	Hidden AI use, uneven skills, or poorly managed role changes erode trust and productivity	Clear permission to use approved tools, training, consultation on significant changes to roles	People and culture
10	Trust and reputation	An AI error in public communications or service delivery damages the trust of clients, donors, residents or customers	Review of public-facing AI content, incident response plan that covers AI	Whole board
11	Missing the opportunity	We fall behind peers or competitors, or keep stretched staff on manual work that AI could safely reduce	AI included in strategy, a funded pilot with measured outcomes, an annual board discussion of value	Whole board
12	Board oversight gap	The board lacks the knowledge, information or reporting rhythm to oversee AI	Director education, AI reporting in the board calendar, a named executive owner	Chair
13	Recording and capture	AI glasses, pendant recorders or meeting notetakers capture clients, residents, staff or visitors without consent, breaching surveillance or privacy law	Recording policy for staff, visitors and contractors, notice and signage, rules for client and care settings	People and culture

How to use it. Table the list at your next audit and risk meeting and ask management to bring back ratings, owners and controls by the following meeting. Risks 1, 2 and 4 are the ones most organisations will rate highest right now. Risk 11 is the one most registers leave out, and it's the one that keeps the conversation from becoming purely defensive.

This list is general information to start a conversation. It isn't legal or risk advice for your organisation.

One Question for Your Next Board Meeting

“What AI is in our building, and how would we know?”

If the answer is a confident list of approved tools, ask how management knows nothing else is in use, including what staff and visitors are wearing. If the answer is a pause, you've found your first agenda item.

If you'd like help running an anonymous AI usage survey or mapping your automated decisions before 10 December, get in touch and we'll set up a conversation. SalesFix can also run education sessions for boards and executives. And if you think I've got something wrong, I'd like to hear that too.



The AI Board Briefing is general information for directors and executives. It isn't legal, financial or risk advice for your organisation.

Sources

- Employment Hero and PagerDuty survey findings: [HRD Australia](#), 7 July 2026
- UK National Cyber Security Centre: [The hidden risks of shadow AI](#), 7 September 2026
- Victorian child protection case: [Parliament of Victoria Hansard](#), 16 October 2024
- Deloitte refund: [CFO Dive](#)
- AI and legal disciplinary matters: [Lawyers Weekly](#), 12 August 2026
- Bunnings tribunal decision and Children's Online Privacy Code: [Bird & Bird AI Regulatory Horizon Tracker](#), Australia
- Automated decision-making obligation: [Gilbert + Tobin](#)
- ASIC review of bank AI use: [Reuters via US News](#), 30 September 2026
- Australian AI policy timeline: [SafeAI-Aus](#)
- Meta Muse launch: [Meta Newsroom](#), 8 September 2026
- Muse training opt-out: [CNBC](#), 8 September 2026
- Amazon blocking Muse: [CNBC](#), 27 September 2026
- Comparison of eight personal agents: [Every](#)
- Work versus personal use of agents: [The AI Daily Brief, How to Choose Your Personal AI Agent](#), 4 October 2026
- Salesforce and YouGov survey of Australian workers: [SecurityBrief Australia](#), 22 June 2026
- Slackbot general availability: [Business Wire](#), 13 January 2026
- Kmart Anko Smart Glasses and workplace recording law: [Hall & Wilcox](#)
- Proposed Commonwealth workplace smart glasses ban: [Reuters via The Edge](#), 17 September 2026
- OAIC investigation of smart glasses: [SmartCompany](#)
- Lawyers on banning smart glasses at work: [Lawyers Weekly](#), 23 September 2026