Expert-Grounded AI

EGAI.COM.AU

Quarterly eNews

NOVEMBER 2026

Q4 2026 | CLIENT EDITION

In this issue

This quarter's regulatory and practice updates, with actions for education and training providers, technology suppliers and the people accountable for AI use.

01	Welcome: put the guidance to work	2
02	The 10 December privacy deadline starts with your decisions	3
03	ASQA brings AI into everyday self-assurance	4
04	Commonwealth AI policy: know which requirements apply	5
05	Workplace AI adoption needs the people who do the work	6
06	Frontier AI rules will not cover every education risk	7
07	Public-sector AI use grows. Are the benefits being tested?	8
08	AI.gov.au: build capability around a real work task	9
09	Updated resources for governance, oversight and readiness	10
10	Your year-end AI adoption agenda	12

THIS QUARTER | Turn updated guidance into operating practice

Privacy transparency, personnel capability and stronger oversight of AI use. Know what applies to your organisation, update the relevant controls and check that they work.

Short on time? Start here

1. Privacy: before 10 December, check whether any automated decisions need to be explained in your privacy policy (Section 02).
2. ASQA: the Version 2 Practice Guides bring AI into training, assessment and learner information (Section 03).
3. EGAI resources: download the current editions from your client link (Section 09).

01 WELCOME

Put the guidance to work

The next step is to show how the guidance changes what people do, what students experience and who remains accountable.

Welcome to our November EGAI update.

In August, we asked organisations to make their AI use visible. Which tools were operating? What were they doing? Who was responsible for checking the result? This quarter, we are taking that work into the decisions, records and everyday practices that sit behind responsible adoption.

The developments in this edition have different consequences. The December privacy deadline calls for a review of actual decision-making processes. ASQA's updated Practice Guides bring AI into training design, assessment and student information. Commonwealth agencies have their own governance requirements, while the draft workplace guide invites employers to work with their people on how AI changes the job. The proposed frontier AI rules raise a further question about the risks that remain inside each organisation. The articles that follow explain the scope and status of each update.

For education and training, the learning question remains close to every technology decision. Does the activity help students develop and demonstrate the skills they need? Can educators recognise a weak or incorrect output? Have we preserved the time, authority and evidence needed for people to exercise judgement?

One pattern keeps showing up in our work with educators. When AI-generated feedback is copied and pasted straight into a student's assessment, the errors tend to travel with it. When an educator types the feedback out, they not only organically rewrite it in their own words, they start noticing what needs to change: the comment that doesn't fit *this* student, the point that isn't quite right, the tone that's simply 'off'. Same tool, very different result. The difference is whether the AI workflow gives a person a moment to genuinely engage with the output. That moment is where professional judgement happens, and it's worth designing in rather than designing out.

Phill and I have updated the EGAI resources to reflect ASQA's Practice Guides v2 and the Australian Government's AI policy v2.1. The changes connect governance, use-case assessment, transparency, supplier assurance and operational readiness. They also keep direct agency requirements separate from measures that other providers may adopt or need to address through a contract.

Thank you for continuing to work through these questions with us. Please use this edition with the people who own the decisions: your education leaders, quality team, privacy lead, technology suppliers and those doing the work. Select the changes that apply, assign responsibility and test one live workflow before relying on an updated document.

All the best,

Deniese
Dr Deniese Cox
Co-Founder, EGAI

A question for your next leadership meeting

Where has AI changed a decision, a learning activity or a person's access to support and what evidence tells us that the change is working?

About this edition

This quarterly update supports organisations on board with EGAI. It brings together source guidance, EGAI interpretation and practical actions. The client resource link appears in the resource-update section. The recommendations should be applied to your organisation's circumstances rather than treated as a universal checklist.

02 PRIVACY AND TRANSPARENCY

The 10 December privacy deadline starts with your decisions

The new transparency obligation can apply even when a person makes the final decision. A general statement that your organisation uses AI will not explain the relevant practices.

From **10 December 2026**, Australian Privacy Principle (APP) entities must include additional information in their privacy policy where the APP 1.7 conditions are met. The OAIC released supporting resources on 30 September, including updated APP 1 Guidelines, fact sheets and a decision flowchart.[1]

The trigger has three parts: the entity has arranged for a computer program to make a decision or do something substantially and directly related to making it; the decision could reasonably be expected to significantly affect an individual's rights or interests; and the program uses personal information about that individual. This is a test of how the process works, not whether a supplier calls its product AI.[1,2]

The OAIC's education example is instructive. A private tertiary institute uses software to rank scholarship applications. A human panel makes the final decision and can disregard the ranking, but the recommendation is still substantially and directly related to that decision. Human involvement does not automatically remove the transparency obligation.[2]

Where the conditions apply, the privacy policy must identify the kinds of personal information used, the kinds of decisions made solely by the program and the kinds of decisions supported by the relevant automated activity. The obligation concerns computer programs more broadly than generative AI.[1,2]

For providers, we recommend reviewing admissions, scholarships, learner-risk flags, access to support and employment-related decisions. These are screening priorities, not a claim that every use is captured. Record what the system contributes, which personal information it uses and how significantly the decision could affect the person. Confirm the organisation's APP coverage and any relevant exemptions before settling the disclosure.

EGAI view: discovery comes before disclosure

A privacy policy cannot accurately describe a workflow that the organisation has not mapped. Start with the decision and the information flow, then write the explanation.

Learning design view: learners want to know what happens to their information

Disclosure isn't only a legal task. For learners it starts before and during enrolment, and they bring real questions with them. In our work with adult learners, we've observed many assume a digital skills test is there to screen them out. Learners with an accessibility or adjustment need often don't know whether to mention it, who will see it or what it will change.

A privacy policy won't answer those questions at the moment they truly matter. Where an automated process uses this kind of information, tell learners in plain language, at the point you ask the question: what it's for, who sees it and how it genuinely helps the learner.

Recommended actions

Focus	Put it into practice
Find the decisions	Review both AI and conventional automation, including recommendations produced inside existing systems.
Test the scope	Document the APP 1.7 assessment and obtain privacy advice where coverage or the decision's significance is uncertain.
Publish and maintain	Update the relevant policy disclosures before 10 December and make future workflow changes a review trigger.

Sources: 1. OAIC: new transparency resources, 30 September 2026 | 2. OAIC: APP 1 Guidelines v2.0, including Example G

Read the OAIC transparency resources

03 VET PRACTICE GUIDANCE

ASQA brings AI into everyday self-assurance

Version 2 moves the discussion into the practical evidence behind training, assessment and student information.

ASQA released its **Version 2 Practice Guides on 29 October 2026**, with more AI-related examples, risks, case studies and answers to provider questions. They explain regulatory expectations under the revised Standards for RTOs. They are not a separate amendment to the Standards or a prescribed set of processes that every provider must copy.[3,7]

In training, the guidance asks providers to justify the use of AI-supported learning and show how it sits alongside human instruction, practice and feedback. AI activities should supplement, rather than replace, hands-on practice where the training product requires physical or practical skills. Providers are also asked to assure the accuracy of AI-generated material and evaluate whether the approach improves engagement and outcomes.[4]

Assessment guidance connects AI use to the principles of assessment and rules of evidence. Instructions should explain whether AI is permitted and under what conditions. Assessors need evidence of the student's own skills and knowledge, including explanation or practical demonstration where needed. The guide also identifies risks from AI-assisted marking without clear oversight and from excluding AI-related assessments from validation samples.[5]

The Information Practice Guide brings this work forward to enrolment. Relevant information should explain AI use in training or assessment and implications such as technology requirements and privacy. It also identifies inadequate human oversight of AI-produced marketing and enrolment information as a risk.[6]

EGAI's recommendation is to trace these examples through the existing quality system. A statement in an AI policy should agree with the assessment instructions, the learner information, the trainer's practice and the evidence retained. Review one actual assessment or delivery sequence to see whether those parts tell the same story.

Keep the student's competence visible

A polished response can help present evidence, but it cannot replace the required demonstration of knowledge and skills. Decide what the assessor needs to see, hear or observe in that particular task.

Recommended actions

Focus	Put it into practice
Training and information	Check one AI-supported delivery activity and its learner information against the training product and cohort needs.
Assessment and validation	Check permitted-use instructions, assessor evidence and the inclusion of relevant AI-related risks in validation.
Quality records	Record any internal non-conformance, assign the correction and verify the change in practice.

Sources: 3. ASQA: Version 2 Practice Guides, 29 October 2026 | 4. ASQA: Training Practice Guide | 5. ASQA: Assessment Practice Guide | 6. ASQA: Information Practice Guide | 7. ASQA: About the Practice Guides

Explore ASQA's updated Practice Guides

04 AUSTRALIAN GOVERNMENT POLICY

Commonwealth AI policy: know which requirements apply

The updated policy provides a useful governance reference, but its direct agency obligations should not be applied indiscriminately to every education provider.

The Digital Transformation Agency published **AI policy v2.1 on 21 October 2026**, describing the changes as minor updates to the v2.0 framework that took effect on 15 December 2025. It applies to non-corporate Commonwealth entities, subject to the policy's exceptions. The full governance framework was not introduced for the first time in October.[8]

The policy connects agency accountability with individual use cases. It requires accountable officials, owners for in-scope uses, internal registers and public transparency statements. Agencies must share their registers with the DTA every six months. Responsible-use arrangements also include personnel training and processes for reporting and addressing AI safety concerns and incidents.[9,10]

Impact assessment begins during design. An in-scope use case must have its assessment completed and agreed risk treatments applied before deployment. The policy distinguishes the commencement of assessments for new uses from the deadline for bringing previously unassessed existing uses into the required arrangements.[11]

Implementation dates to distinguish

Date	Agency milestone
15 December 2026	The 12-month milestones include use-case owners, an internal register, responsible-use arrangements and mandatory personnel training. Agencies must begin use-case assessments by this point.[8,9,10,11]
30 April 2027	For existing uses not yet assessed, agencies must determine scope and apply the relevant policy actions.[11]

For other education and training providers, separate three questions. Is your organisation directly within scope? Has an agency contract imposed relevant requirements? Which measures would improve your own governance even where they are not mandatory? Those answers determine what belongs in your policy, approval record and supplier arrangements.

A government client may need evidence about your AI components, data handling or safeguards to discharge its own responsibilities. Check the actual contract and agency instructions rather than assuming that all agency reporting duties transfer to the provider. Similarly, the policy's recommended technical and procurement guidance should not be relabelled as universally mandatory.[10]

EGAI view: record why a requirement applies

Mark each requirement as a direct obligation, a contractual condition or an adopted governance measure. This keeps approvals and assurance proportionate and prevents an agency policy from being mistaken for a general RTO requirement.

Sources: 8. DTA: Policy for the responsible use of AI in government v2.1 | 9. DTA: Strategy and oversight | 10. DTA: Preparedness and operations | 11. DTA: AI use case impact assessment

Read the policy's strategy and oversight requirements

05 WORKFORCE CONSULTATION

Workplace AI adoption needs the people who do the work

The new draft workplace guide brings work design, worker engagement and capability into the adoption decision.

On **12 November 2026**, the Australian Government released the draft **Best Practice Guide on the Adoption of AI in the Workplace**. Consultation closes at **11:59 pm AEDT on 27 January 2027**. It is a draft best-practice guide for consultation, not a new binding workplace standard.[12,13]

Organised around Plan, Enable and Embed, the draft sets out eight principles spanning opportunities, readiness, worker engagement, job design, capability, workforce change, governance and work health and safety. It has been developed through engagement with employers, unions, small business representatives and other contributors.[12,13]

For education providers, the practical value is the attention it gives to the work itself. Before introducing AI-generated learner feedback, for example, ask how an educator will check it, how much time that review will take and how students will reach a person when the feedback is wrong. The workflow needs to support the educator's judgement rather than simply allocate more checking to an already full workload.

The same applies to administration. A tool may reduce drafting time while creating new tasks in verification, privacy review or exception handling. We recommend comparing the complete workflow before and after the change, including rework and the experience of the people using the service. A quicker first draft is only one part of that comparison.

Worker involvement should occur while the use case is being designed. Trainers, assessors, student-support personnel and administrative personnel can identify the decisions that need context, the information that is often missing and the circumstances in which automation would make a task harder. Bring those observations into procurement, approval and testing.

EGAI view: build capability through the work

Give personnel a relevant task, clear boundaries and time to practise checking the result. Confidence with a tool should be accompanied by evidence that the person can recognise an error, explain a decision and escalate a concern.

Recommended actions

Focus	Put it into practice
Discuss a live use case	Ask affected personnel where AI would help, where it would add work and which decisions need to remain with people.
Test the whole workflow	Compare quality, review time, rework and user experience before deciding whether to expand the use.
Contribute to the draft	Use an education or training example to explain what the guide should clarify before consultation closes.

Sources: 12. DEWR: draft workplace AI guide announcement, 12 November 2026 | 13. DEWR: consultation announcement and closing time

Read the draft and contribute to the consultation

06 AI SAFETY AND ACCOUNTABILITY

Frontier AI rules will not cover every education risk

The proposed national protections deserve attention. Providers still need to examine the everyday decisions made with smaller and embedded systems.

On **10 November 2026**, the Australian Government announced proposed mandatory protections for frontier AI companies operating in Australia. The six areas cover security, serious-risk management, independent model evaluation, incident notification, published safety frameworks and responses to evolving standards. The Government also proposed a targeted AI safety regulator. Detailed legislation and implementation arrangements remain to follow.[14]

An analysis in The Conversation questions whether the proposed focus on frontier models adequately addresses harms from smaller AI systems embedded in everyday services.[15]

Consider a learner-support system that ranks who receives a call, a tool that drafts assessment feedback or an admissions process that filters applications. A technically secure model could still be used in a workflow that misdirects support, produces misleading feedback or removes a professional review.

EGAI's view is that risk assessment must make the possible reduction of human judgement visible. Which decisions will people stop making? What will they still understand about the evidence? Can an authorised person question the recommendation, change the outcome or stop the process? A final approval button does not answer those questions.

This is also a reason to test ordinary failure, not just a major security event. Try incomplete information, an unusual learner circumstance, a conflicting source or a request outside the tool's purpose. Examine what the system does and whether personnel can recognise and correct the result.

Assess the use, not the prestige of the model

A supplier's safety controls can inform procurement. They do not establish that a specific assessment, support or employment workflow is suitable for your organisation.

Recommended actions

Focus	Put it into practice
Separate the obligations	Track the proposed frontier-company rules without presenting them as commenced duties for every AI user.
Review local decisions	Identify affected people, potential harms, human authority and a usable route to challenge an outcome.
Test intervention	Demonstrate that personnel can pause the workflow, correct the result and restore the service when something goes wrong.

Sources: 14. Australian Government: AI safety announcement, 10 November 2026 | 15. The Conversation, 11 November 2026: Australia will impose mandatory safety rules on AI. Are we focused on the right risks?

Read The Conversation analysis

07 EVIDENCE FROM PRACTICE

Public-sector AI use grows. Are the benefits being tested?

The Mandarin and Liquid's research offers a useful comparison for providers: access and regular use do not establish the quality of the result.

The **2026 AI and the Australian Public Service Special Report**, published by The Mandarin on 3 November, draws on the Let's Get Real About AI survey conducted with Liquid. Its findings cover **694 respondents** from federal and state or territory government. They should not be treated as a census of the APS or as estimates for the education sector.[16,17]

Among respondents, **94% had access to AI tools and 83% used them at work**. Yet **74% said outputs needed some or significant correction**, while only **8% reported wide organisational deployment**. The report describes much of the use as individual and task-based, particularly drafting, analysis and routine administration.[16]

For education leaders, these findings suggest a better set of questions than how many accounts have been activated. What work is being produced? Who checks it? How often does it need correction? Does the revised process save time after review and rework are included? These are questions for local testing, rather than conclusions the survey can answer for a provider.

We recommend taking one frequently used activity through a small, documented evaluation. For an AI-assisted learning resource, retain the original task brief, check training-product alignment and have a competent educator review factual accuracy and teaching suitability. Record the changes needed and the total time taken. Compare that with the existing method before making a claim about efficiency.

Keep the learner's experience in the evaluation. A shorter preparation process is of limited value if the resource becomes less clear or the educator has less time to respond to questions. Similarly, a support chatbot should be assessed on the usefulness of the help and the route to human assistance, not simply the number of conversations completed.

We're also seeing more "quick bots" in practice: chatbots added to a course or support page without the guardrails to keep them contained. If a bot isn't confined to expert-grounded, pre-approved sources, it can drift outside its purpose or answer from general content that doesn't match the organisation's own policies or training products. Other bots can quote a source but not reason with it, so users get a passage of text rather than help with their actual question. Before a bot goes live, check what it is allowed to draw on, what it does when a question falls outside that boundary and whether it helps people apply the information, not just find it.

EGAI view: use measures that can change a decision

Agree the quality, time and student-impact measures before the trial. Record what would justify continuing, changing or stopping the use case.

Recommended actions

Focus	Put it into practice
Choose a baseline	Document the existing method, review effort and service or learning outcome before testing the alternative.
Count the checking	Include correction, verification and escalation in the evaluation, not only time spent generating an output.
Decide from evidence	Use the result to retain, amend or stop the workflow. Treat confidence and satisfaction as inputs, not proof of capability.

Sources: 16. The Mandarin: Let's Get Real About AI findings, 3 November 2026 | 17. The Mandarin: 2026 AI and the Australian Public Service Special Report

Explore the 2026 public-sector AI special report

08 CAPABILITY AND FREE LEARNING

AI.gov.au: build capability around a real work task

The National AI Centre's resources can support a personnel development plan when learning is connected to the work people are authorised to do.

AI.gov.au brings together practical guides, learning resources and event listings. Its training directory covers introductory AI, generative AI, responsible use, privacy and agentic AI. It also links to free AI and emerging-technology modules through the Australian Government's Digital Solutions Program. Costs and eligibility vary across the wider directory, so check the individual offering. [18]

A recent addition is **Understanding and using agentic AI**, published on 26 October. The resource explains systems that can plan steps, use tools and take actions towards a goal. It covers potential uses, risks, limits, approvals, oversight and implementation. This is relevant when a proposed tool can act in a system rather than only provide a draft for someone to review. [19]

For providers, we recommend choosing learning according to a person's responsibilities. An educator reviewing feedback needs practice checking evidence and explaining permitted student use. A procurement officer needs to examine data access and supplier claims. A manager approving an AI agent needs to understand what it can do, when permission is required and how its activity can be interrupted.

Apply the learning at work

Use these learning resources as a starting point for a workplace activity. Ask each participant to bring back one proposed change, identify the approval or review it needs and demonstrate it on a safe example. Attendance records show participation; a checked work sample helps show how the learning is being applied.

Keep the organisation's boundaries in the room

A public course or demonstration does not approve a tool for your organisation. Use approved environments, avoid introducing personal or confidential information into exercises and test the method before using it with students or clients.

Sources: 18. National AI Centre: training and self-directed learning | 19. National AI Centre: Understanding and using agentic AI, 26 October 2026

Browse AI.gov.au learning resources

Read Understanding and using agentic AI

09 EGAI RESOURCE UPDATE

Updated resources for governance, oversight and readiness

The latest release follows our ASQA Practice Guides v2 update and adds the Commonwealth agency requirements reflected in AI policy v2.1.

We have completed a further update to the **EGAI AI Adoption resources** to reflect the Australian Government's updated AI policy for government entities. The changes connect the source requirements with the strategy, registers, assessments and records used to govern actual AI activity.

The resources distinguish requirements that apply directly to Commonwealth agencies from governance measures education and training providers may adopt. They also identify where agency requirements may become relevant through a contract. This distinction is carried into the operational resources so that users can record the basis for a requirement, rather than assume it applies to every organisation.

What the update covers

Area	What has been addressed
Governance and strategy	Accountable officials, use-case ownership, personnel capability, incident management and reporting obligations.
Register and transparency	Use-case registration, AI classifications, transparency statements and reporting requirements.
Assessment and approval	Impact assessment, risk ratings, approval decisions, safeguards, residual risks and reapproval requirements aligned with DTA policy v2.1.
Supplier due diligence	Supplier assurance, embedded AI components, intellectual property, data handling, contractual protections and exit arrangements.
Agentic AI	Human oversight of systems that can take action, including permitted actions, monitoring, intervention and safe shutdown.
Operational readiness	Controls for moving from pilot activities into broader operational deployment.

These changes sit alongside the recent ASQA Practice Guides v2 update. That earlier work brought the revised education guidance into the EGAI resources, including learner information, evidence of learning and assessment, human review and the checks needed before a pilot expands.

One use case, connected records

The strategy, register, approval record and supplier review should describe the same activity. A change to the data, permitted actions, affected people or level of human review should be visible across the relevant records.

Sources: 8. DTA: Policy for the responsible use of AI in government v2.1 | 3. ASQA: Version 2 Practice Guides, 29 October 2026

09 EGAI RESOURCE UPDATE | CONTINUED

Bring the updated resources into your working system

The updated resources are available now through the **existing EGAI client link**. Download the current editions and check the versions used in shared folders, internal templates and any authorised AI workflow that draws on the Intelligence Packs.

Open the current EGAI client resources

Resource	Current version
EGAI AI Adoption Toolkit	**v1.4**
EGAI Intelligence Packs for AI Adoption	**v1.4**
EGAI Specification Paper: AI Adoption in Australian Education and Training	**v1.5**

Apply the changes to a live use case

Choose a use case already in operation or about to move beyond a pilot. Confirm its owner, check the privacy and supplier information and review the approval conditions. Record any unresolved evidence gap and the decision about whether the use can proceed. Retain the reasons as well as the result.

For teams using the Intelligence Packs, update the controlled files used by the authorised AI system and test a representative task. Check that the output distinguishes agency requirements from provider good practice, identifies missing information and leaves approval decisions with the authorised person. A new file version does not itself verify the workflow.

The resources can support internal implementation and, where the EGAI licence permits, commercial client work. Keep each client's circumstances, data and approval decisions separate. Use the materials to structure the work, not to imply that a product or organisation has received government approval.

WA TAC guidance reviewed

We have also reviewed the recently updated WA Training Accreditation Council fact sheets relevant to WA TAC-regulated VET providers. That review identified **no further amendments needed to the EGAI resources**. This is a finding about the resources, not a statement that providers can disregard the updated guidance or skip a review of their own practices.

A more manageable update cycle

The ASQA and Commonwealth policy updates arrived close together. Where multiple guidance changes are released in quick succession, we will consolidate resource changes into **a maximum of one update per week**. This maintenance continues as regulatory guidance and requirements develop.

Keep the client link within the licensed access arrangements

Use the existing folder for the current editions. Preserve your organisation's approved records and local adaptations when incorporating a new version and retain an identifiable superseded copy where needed for traceability.

10 YOUR NEXT STEPS

Your year-end AI adoption agenda

Use the deadlines to organise the work. Use evidence from actual activities to decide whether the changes are effective.

Not every development in this edition creates the same obligation. The privacy changes have a commencement date, the Commonwealth policy has a defined organisational scope, the workplace guide is open for consultation and the frontier AI protections are proposed. Keep those distinctions in your implementation plan.

Dates to put in the diary

Date	What to address
10 December 2026	**Privacy transparency:** APP entities should have the required privacy-policy information in place for arrangements meeting APP 1.7. [1,2]
15 December 2026	**Commonwealth agency implementation:** the policy's 12-month milestones fall due, including the start of use-case assessments. Check the requirements and scope relevant to the agency. [8,9,10,11]
27 January 2027	**Workplace AI guide:** consultation closes at 11:59 pm AEDT. [13]
30 April 2027	**Existing agency uses:** complete scope determinations and relevant policy actions for existing uses not yet assessed. [11]

Three priorities before the next issue

Priority	Evidence to bring back to leadership
1. Make a decision visible	Select a student, workforce or client-facing decision influenced by AI or automation. Map the information, system contribution, human authority and review pathway.
2. Update the working records	Incorporate the current EGAI resources where relevant. Confirm that the register, supplier review, approval conditions and disclosures agree.
3. Test the result	Check one real workflow against its expected benefit and safeguards. Record quality, rework, affected-person experience and any change to human judgement.

Need an independent read on a live use case?

Bring the purpose, proposed workflow, data involved and the decision you need to make. Deniese and Phill can help you work through the regulatory and learning-design questions using your EGAI resources.

Contact Deniese and Phill

Sources: 1. OAIC: new transparency resources, 30 September 2026 | 8. DTA: Policy for the responsible use of AI in government v2.1 | 11. DTA: AI use case impact assessment | 13. DEWR: consultation announcement and closing time