

B1 — Briefing for Deputy and Pro Vice-Chancellors, Deputy Provosts, and Vice-Presidents of Research

The five decisions about AI in doctoral research that sit with you and nobody else, the five questions to ask your people this month, and what it costs to find out where [INSTITUTION NAME] actually stands.

Browse the full online instrument catalogue to read or download every resource in HTML, Word, PDF, and spreadsheet formats.

About this instrument

Purpose	To decide, in one sitting, whether to commission an institutional AI-readiness assessment, based on evidence rather than general impressions.
Who reads it	Deputy Vice-Chancellors and Deputy Provosts of Research, as well as Pro Vice-Chancellors and Deputy Vice-Presidents of Research. There is nothing to complete.
Time required	About 10 minutes to read, and one hour to act on the "Where to start" section.
Related report sections	Part 4 and Appendix A (the maturity grid), §5.1 (executive research leadership), §1.3 (the institutional baseline), and §2.5 (Dimension 5 and the four axes).
Cells touched	10 of the framework's 20 cells (five dimensions across four axes: policy, people, systems, and process). All four cells of Dimension 5 (institutional benchmarking grid), plus D1-policy (human-in-the-loop discipline), D2-policy (responsible use in practice), D3-policy and D3-systems (tooling that promotes responsible use), and D4-people and D4-process (AI-literate humans).

Table 1. Summary of the purpose and operational specifications of the institutional assessment instrument.

Why this is on your desk now

The companion report, *Responsible AI in Academic Research: A Competency Framework for Research Training*, scores institutions across five dimensions at four developmental levels: *absent*, *nascent*, *established*, and *leading*. On the fifth dimension, which covers the institutional benchmarking grid, not a single university in the report's sample of 38 top-tier doctoral institutions across 15 countries and jurisdictions publishes its own self-assessment against the framework. For this reason, the report's §4.2 describes the *leading* level as "a forward-looking standard that the framework's adopters can elect to populate." In this framework, developmental progress is organized like steps on a ladder, where each improvement is a single rung. In other words, the report names a rung that nobody currently occupies, and this pack provides the practical instruments to help your institution reach it.

- **Practice has outrun disclosure.** When the *BMJ* added a structured disclosure field to its manuscript submission forms, only 5.7 percent of 25,114 submissions across 49 *BMJ* journals disclosed any AI use in the seven months that followed (§1.1). In reality, AI assistance is already deeply embedded in literature reviews, code development, drafting, and increasingly statistical reasoning. This contrast suggests that what graduate students and researchers are actually doing in practice has run far ahead of formal disclosure.

- **The usual framing puts the problem in the wrong office.** Fifteen of the 38 universities in the study (about two in five) currently maintain AI policies whose substantive scope stops at plagiarism (§1.3, §3.2). But in doctoral research, this is fundamentally a research integrity question involving scientific validity, reproducibility, citation verification, and adversarial review, rather than an academic integrity question about who typed the words.
- **The leading edge is thin.** Only six of the 38 universities reach past baseline research integrity into active AI literacy and valid research practices (§1.3, §3.2). Academic publishers converged on a common rule prohibiting AI authors within about three months (§1.1). In contrast, universities have not yet converged on a shared standard, and institutions that establish a clear position now will help shape the sector-wide standard that emerges.

For Australian readers. TEQSA's *Gen AI strategies for research training: Emerging practice* (June 2025) sets out what institutions should do, organizing recommendations into four sections of checklist actions alongside case studies from named Australian institutions. The Process / People / Practice triad comes from its companion report, *Gen AI strategies for Australian higher education: Emerging practice* (November 2024), which is the closest existing institutional readiness framework to Dimension 5. Dimension 5 builds directly on that work by adding an explicit policy axis and making a clear distinction between systems and processes (§3.4). While TEQSA sets out what universities should do, this pack supplies the concrete instruments for putting those recommendations into practice and measuring how well your institution is doing.

The five decisions only you can make

1. Whether AI in doctoral research is treated as a research integrity matter or a student conduct matter. This decision sets which office drafts institutional policy, which committee adjudicates alleged breaches, and what standard of evidence applies. At most institutions, this was initially settled by default in favor of student conduct or academic integrity. Reversing that default across the Registrar, the Research Integrity Office, and the Graduate School requires your executive authority. (*D1-policy, D2-policy*)

2. Who owns the institution's benchmarking cycle, and how often it runs. This requires one named person with a published mandate and a clearly stated reporting schedule. The role can sit with a Director of Research Integrity, a Dean of Graduate Studies, or the chair of [COMMITTEE]. However, the appointment itself cannot be delegated down, because the owner must have the standing and authority to require documentary evidence from IT, the library, the graduate school, and every school and college across the university (§5.1). (*D5-people, D5-process*)

3. Whether one institutional gate governs which AI tools may touch unpublished research data. At present, the default at most institutions is that tool selection is left to individual researchers and laboratories. Only you have the authority to establish a common review gate across all schools and colleges, and you are the executive leader who will be asked to explain what happened if unpublished research data leaves the university through an unvetted consumer AI account. (*D3-policy, D3-systems*)

4. What is funded, both upfront and on a recurring basis. Dimension 3 (*tooling that promotes responsible use*) carries the largest upfront cost, which includes developing a procurement standard, conducting sampled testing of vendor outputs, and completing rights impact assessments. Dimension 4 (*AI-literate humans*) carries the largest recurring cost, because curriculum development, supervisor training, and examiner briefings must be repeated every year as tools and regulatory rules change (§5.1). An institutional policy that lacks a dedicated recurring budget line is really just a published statement of good intentions rather than a functioning capability. (*D3-systems, D4-people*)

5. What the institution states publicly, and who is named accountable if a regulator asks.

Publishing your readiness scores, participating in external or peer audits, and contributing to sector-wide standards supply evidence toward D5-policy → leading, D5-people → leading, and D5-process → leading, respectively (Appendix A). A fresh A1 assessment determines whether all four Dimension 5 cells qualify at *leading*. Deciding what information to make public is a key reputational judgment. At the same time, designating a specific accountable officer for the regulatory obligations that apply in your jurisdictions is a critical legal compliance decision. (D5-policy, D5-people, D5-systems, D5-process)

Five questions to ask your people this month

Each of these five questions can be answered directly by an existing document or artifact. If the document does not exist, that absence is itself the finding. I recommend looking past the standard high-level reassurances. For instance, the statement that "AI-generated content must be verified and should not replace expert opinion or judgement" is taken verbatim from the Australian Research Council and NHMRC guidance (§1.2), but it does not give researchers or students any practical instructions on what to actually do.

Question to ask (and of whom)	An unhelpful answer	The artifact that settles it
Which document tells a researcher or graduate student what to do when an AI tool returns a plausible-looking citation, and which clause? (Ask the Director of Research Integrity)	A restatement of general principles without pointing to a specific clause.	A document with a named clause that provides an actionable task-level rule, its last review date, and details on the training module where it is taught.
Show me the last five completed thesis or dissertation declarations. What did the graduate student declare about AI, and what were the examiners (dissertation committee members, in US usage) instructed to check? (Ask the Dean of Graduate Studies)	"The form has a check box."	Five completed declarations, plus written examiner instructions specifying what to verify during the oral examination (the viva or thesis defense).
Which AI tools are approved to process unpublished research data, who approved them, and against what test? (Ask the Chief Information Officer and Data Protection Officer)	"We hold an enterprise license" (which is a procurement contract, not an operational gate).	A published procurement standard, a named approver, and dated records of vendor outputs that were sampled and tested at deployment.
Which doctoral training course carries citation verification and model-and-parameter specification as named learning outcomes, and how many graduate students completed it last year? (Ask the Dean of Graduate Studies and the methods curriculum owner)	"It is embedded across the training program."	A course outline or syllabus listing those learning outcomes by name, along with verified completion numbers for the last full academic year.

Question to ask (and of whom)	An unhelpful answer	The artifact that settles it
Which of our AI systems fall inside the regulatory classifications that apply to us, who is accountable for each assessment, and by what date? (Ask the General Counsel and Data Protection Officer)	"We are monitoring developments."	A register of in-scope systems, a named accountable officer for each, and scheduled completion dates that precede the regulatory deadline.

Table 2. Key questions for institutional personnel, examples of unhelpful answers, and verifying artifacts.

On the fifth question. The regulatory classifications that most commonly apply to universities cover admissions, the evaluation of learning outcomes, the assignment of education levels, and exam proctoring. Two compliance dates are already fixed (§2.5). The first is the **EU AI Act Article 27** fundamental-rights impact assessment obligation, which takes effect on **August 2, 2026**, for institutions operating in or serving the EU market. The second is the **Australian Privacy Act 1988** automated decision-making transparency obligation, which takes effect on **December 10, 2026**, for institutions operating under the Australian privacy regime. Other jurisdictions carry equivalent requirements, so the core question remains the same everywhere.

The 90-day sequence, and what it costs

Days	What happens, and who does it	Effort	Evidence supplied
1–15	Designate the owner of the benchmarking cycle, and publish their mandate and reporting schedule. You can accomplish this in a single executive meeting.	1 hour	Evidence toward D5-people → nascent. Dimension 5 advances only after a fresh A1 rescore of all four D5 cells.
16–45	Collect documentary evidence using the maturity self-assessment (A1). The owner and an analyst request documents from research integrity, the graduate school, the library, IT security, and data protection. A graduate school administrator fields the graduate student and advisor/supervisor surveys (A2, A3). You or the owner brief deans, graduate school leadership, the research integrity and graduate studies committees, and the graduate student representative body (B2–B5).	About 10 person-days, plus 2 to 10 minutes per survey respondent	An evidence base for all 20 cells
46–60	Score the 20 framework cells in a single workshop using A1 , and identify the binding axis for each dimension. This brings together 8 to 12 colleagues across research integrity, the graduate school, the library, IT security, data protection, two Associate Deans of Research, and one graduate student representative.	Half a day (about 5 person-days total)	A defensible dimensional profile

Days	What happens, and who does it	Effort	Evidence supplied
61–75	Complete the gap analysis and develop the action ladder (A4), choosing specific improvements to make over the next 12 months, each with a named owner and deadline. This is done by a smaller working group from the workshop.	Half a day (about 3 person-days total)	Named, dated commitments
76–90	Prepare the board reporting one-pager (A5) to communicate the institutional profile and readiness status for applicable regulatory obligations. The owner drafts the report, and you review and sign off.	About 2 person-days	Evidence toward D5–process → established after formal tabling on the required schedule. Dimension 5 advances only after a fresh A1 rescore of all four D5 cells.

Table 3. Implementation schedule showing required actions, estimated effort, and expected progress over ninety days.

This represents roughly 22 person-days of effort across one quarter, plus about ten minutes from each graduate student who completes the survey. It requires no new IT systems, no vendor procurement, and no external spending, unless you elect to participate in external benchmarking, which supplies evidence toward D5–people → leading (§5.1). If you need an immediate pulse check, a short audit of four to six questions per dimension can be completed in about an hour (Appendix A). That short audit will give you an initial indication of where you sit, whereas the full 90-day sequence produces the rigorous evidence base that a university governing body will accept.

One essential rule about the result. The assessment does not reduce the institution to a single uninformative score, and any tool that attempts to emit a single score is being misused. The output is a clear profile across five distinct dimensions. The report requires all four axes (policy, people, systems, and process) to operate for the *established* level. The pack applies that rule by setting each dimension at the level of its lowest-scoring axis. It calls that limiting axis the *binding axis*. In other words, the true diagnostic value of the assessment lies in identifying the pattern across dimensions and knowing where to focus your effort, rather than chasing a single summary number (§4.3).

Where to start

Name the owner of the benchmarking cycle and set [DATE] as the deadline for their first dimensional score. Send the five questions above to the five colleagues named beside them, and ask for written answers with the supporting artifacts attached. If the artifacts arrive, you will already hold the core of your institutional evidence base. If they do not arrive, you already have your answer about where [INSTITUTION NAME] currently stands.

Drafting note. Replace [INSTITUTION NAME], [COMMITTEE], and [DATE] before circulating this briefing. Substitute local role titles (such as Deputy Provost, Vice-President Research, or Pro-Rector for Research) and the appropriate local term for doctoral researchers, such as HDR (higher degree by research) candidate in Australia, PhD student in the United States, or research student in the United Kingdom. Keep the regulatory dates that apply to your specific jurisdictions, and delete the rest.

Crosswalk: Briefing item to report section

Item in this briefing	Report section
Leading on Dimension 5 unreached in the sample	§4.2
BMJ disclosure rate, and practice outrunning disclosure	§1.1
Two in five of the 38 universities stop at the plagiarism ceiling	§1.3 and §3.2
Research integrity rather than academic integrity framing	The report finding and §1.2
Six of the 38 universities reach AI literacy and valid research practices	§1.3 and §3.2
Publisher convergence on AI authorship within about three months	§1.1 and §5.1
TEQSA Process / People / Practice triad as antecedent to Dimension 5	§3.4
Decision 1: Research integrity framing and the owning office	The report finding, §1.3, and §2.1
Decision 2: Owner and reporting cadence for the readiness score	§2.5 and §5.1
Decision 3: Institutional gate for tools touching unpublished data	§2.3 and §5.1
Decision 4: Upfront costs on Dimension 3, recurring costs on Dimension 4	§5.1, §2.3, and §2.4
Decision 5: Public position and named regulatory accountability	§2.5, §5.1, and Appendix A
Question 1: Task-level rules that researchers can act on	§1.2 and §2.1
Question 2: Thesis declarations and examiner instructions	§2.1, §4.2, and §5.3
Question 3: Approved tools, named approver, and sampled outputs	§2.3
Question 4: Named curriculum outcomes and completion counts	§2.4 and §4.2
Question 5: Register of in-scope systems and accountable officers	§2.3 and §2.5
Regulatory compliance dates: EU AI Act Article 27 and Australian Privacy Act 1988	§2.5 and §2.3
Assessment cost and duration, and the one-hour short audit	§4.1 and Appendix A
No single score, and this pack's weakest-axis rule	§4.1, §4.3, and Appendix A
The 90-day sequence	§5.1 and Part 4

Table 4. Mapping of briefing topics to their corresponding sections within the report.



Michael J. Zyphur, PhD · Professor and Director, Instats · instats.org · support@instats.org

Cite the pack. Zyphur, M. J. (2026). *The Institutional AI Readiness Pack: Self-Assessment and Implementation Tools for Responsible AI in Academic Research*. Instats Policy Series. <https://doi.org/10.61700/bv2nulyhht>

Companion report. Zyphur, M. J. (2026). *Responsible AI in Academic Research: A Competency Framework for Research Training*. Instats Policy Series. <https://doi.org/10.61700/t31oy23grr>

License. The pack and its instruments are licensed under [Creative Commons Attribution 4.0 International \(CC BY 4.0\)](#). You may adapt them for institutional use with attribution.