EW-AiRM™

BY HUMAN-AI.INSTITUTE

FORMAL SUBMISSION

Artificial Intelligence and Data Centres

Submission to the Senate Environment and Communications References Committee, Parliament of Australia, applying EW-AiRM™ (Enterprise-Wide AI Risk Management) to the Inquiry's Terms of Reference

SUBMITTED BY

Prof. Markus Krebsz
Creator & Author of EW-AiRM™
Founding Director,
Human-Ai.Institute

DATE

24 August 2026
Submissions close
1 September 2026

INQUIRY

Referred 13 May 2026
Reporting date
16 November 2026

ONLINE

www.ewairm.com
www.human-ai.institute

Prof. Markus Krebsz
Creator & Author, EW-AiRM™ (Enterprise-Wide AI Risk Management)
Founding Director, Human-Ai.Institute · UNECE WP.6 AI Project Lead
contact@human-ai.institute · www.ewairm.com

Committee Secretary
Senate Standing Committees on Environment and Communications
PO Box 6100, Parliament House
Canberra ACT 2600, Australia
By email: ec.sen@aph.gov.au / via Submission web-portal

24 August 2026

Re: Inquiry into Artificial Intelligence and Data Centres, referred 13 May 2026

Dear Committee Secretary,

I welcome the opportunity to make a submission to the Senate Environment and Communications References Committee's Inquiry into Artificial Intelligence and Data Centres, and I thank the Committee for extending the submission period to 1 September 2026.

I write as the creator and author of EW-AiRM™ (Enterprise-Wide AI Risk Management), a three-layer AI risk governance framework published in full by Wiley Finance in November 2026, and as Founding Director of the Human-Ai.Institute. My work sits at the intersection of the two questions the Committee has been asked to examine: whether existing regulatory frameworks are effective in governing the growth of AI infrastructure, and how the risks that AI and data centres pose to communities, industries, the environment, water and energy can be identified, measured and managed before they crystallise.

The central proposition of this submission is that the Inquiry's Terms of Reference describe, at their core, a risk governance problem. Australia's current settings regulate the inputs to data centres (land, energy, water) and offer voluntary guidance on the AI that runs inside them, but no instrument yet requires the organisations that build and operate this infrastructure to manage AI risk on an enterprise-wide basis, with named accountability, structured pre-deployment assessment, measurable thresholds and protected incident reporting. EW-AiRM™ provides an implementable, internationally aligned blueprint for closing that gap, and this submission sets out how its architecture, its Five Non-Negotiables and its measurement discipline can inform the Committee's findings under each Term of Reference.

I consent to this submission being published in full and will be publishing it also via www.Human-Ai.Institute and www.ewairm.com. I would be pleased to expand on any aspect of it, in writing or by giving evidence to the Committee via videoconference from the United Kingdom, should that assist the Committee's deliberations.

Yours faithfully,

Prof. Markus Krebsz

UN ECE	AI Project Lead (Embedded AI Systems)
The Human AI Institute	Founding Director
Enterprise-wide AI Risk Management (EW-AiRM)	Creator & Author
Stirling University (UK)	Honorary Professor
Woxsen University (India)	Clinical Professor of Practice

unece.org/trade/wp6/digital-regulation-goods-artificial-intelligence
unece.org/trade/documents/2025/03/informal-documents/markus-krebsz-wp6-project-leader-ai
www.Human-Ai.Institute · www.ewairm.com · www.stir.ac.uk/expert/in/artificial-intelligence · www.krebsz.net

contact@human-ai.institute · krebsz@web.de · +44 (0) 79 85 065 045

SECTION 1

About the Submitter and About EW-AiRM™

1.1 The submitter

Prof. Markus Krebsz is the creator and author of EW-AiRM™ (Enterprise-Wide AI Risk Management) and the Founding Director of the Human-Ai.Institute and of De-Risking Solutions Ltd. He serves as AI Project Lead for the United Nations Economic Commission for Europe (UNECE) Working Party 6, where he led the development of the Common Regulatory Arrangement (CRA) on products/services with embedded AI systems (ECE/TRADE/486), adopted in 2024 and available to UN member states for formal signature. He is a member of the European AI Office (DG CNECT) expert group, a plenary participant in the General-Purpose AI Code of Practice process and a member of all four of its working groups, a StandICT.eu Standards Maker contributing to IEEE P7007.1, a ForHumanity Certified Auditor, an Honorary Professor at the University of Stirling (UK) and a Clinical Professor of Practice at Woxsen University (India). He has served since 2017 as an independent non-executive director on the Supervisory Council of a European licensed bank, giving him more than three decades of practitioner grounding in regulated-industry risk governance.

The complete EW-AiRM™ framework is published by Wiley Finance as *Enterprise-Wide AI Risk Management* in November 2026, with a foreword contributed from the International Monetary Fund and endorsements from central bankers, regulators, standards bodies and academics worldwide. Framework resources, tools and reference material are freely accessible at www.ewairm.com.

1.2 What EW-AiRM™ is

EW-AiRM™ is a comprehensive framework for governing artificial intelligence risk across an entire organisation rather than system by system or department by department. It is deliberately technology-neutral, jurisdiction-agnostic and interoperable with the legislative, regulatory and/or standardisation instruments the Committee will encounter throughout this Inquiry, including ISO/IEC 42001, the NIST AI Risk Management Framework, the EU AI Act and Australia's own Guidance for AI Adoption. It is structured in three integrated layers:

- **Strategic layer.** Six pillars covering Strategic Alignment and Necessity Assessment, Organisational Readiness, Data and Technological Maturity, Risk Tolerance and Prioritisation, Governance and Accountability, and Through-the-Lifecycle Monitoring and Adaptability. The first pillar insists that the opening governance question is not how to deploy AI but whether to deploy it at all, a question with obvious application to infrastructure siting and approval decisions.
- **Operational layer.** A systematic mapping of the organisation's exposure against the MIT AI Risk Repository, hosted by MIT FutureTech, which categorises more than 1,700 AI risks across seven domains and twenty-four subdomains. Domain 6 (Socioeconomic and Environmental harms) and its subdomain 6.6 (Environmental harm) speak directly to this Inquiry's second Term of Reference.
- **Resilience layer.** Preparation for low-probability, high-impact "AI Black Swan" events across eight categories, including systemic foundation model concentration and supply chain compromises, through three capabilities: robust foundations, continuous sensing and adaptive response.

Operating across all three layers is HAiPECR™, the framework's ethical overlay, originating in 2023 as the Human-Ai Paradigm for Ethics, Conduct and Risk and listed on the OECD AI Policy Observatory since April 2023. Its seven dimensions (Human oversight and autonomy; Accountability; inclusivity and participation; Protection of data, privacy and neural integrity; Ethics; Conduct risk; and Rights and non-harm including intergenerational sustainability) map to the ten core principles of the UNESCO Recommendation on the Ethics of Artificial Intelligence (2021), to which Australia is a signatory. The final dimension, Rights and non-harm including intergenerational sustainability, treats the environmental footprint of AI as a matter of rights owed to present and future generations, not as a discretionary corporate social responsibility item.

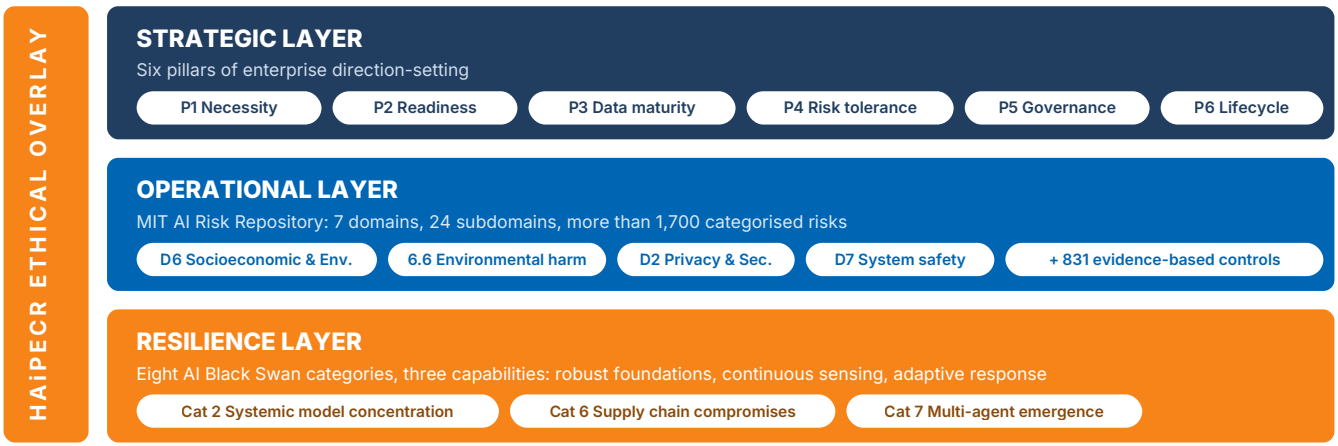

Figure 1. The EW-AiRM™ three-layer architecture with the HAiPECR ethical overlay. Source: Krebsz, *Enterprise-Wide AI Risk Management* (Wiley Finance, 2026); www.ewairm.com. (Displaying only a selection of EW-AiRM™ framework building blocks that are relevant to this consultation)

1.3 Lineage and international standing

EW-AiRM™ is not a rapid response to the current AI cycle. **EW-AiRM™ is the culmination of more than a decade of published governance instruments: the Global Conduct Risk Paradigm (2015), the Universal Conduct Risk Paradigm presented to UNECE in 2017 and published on the UNECE website, HAiPECR (2023, OECD AI Policy Observatory), and the UNECE Common Regulatory Arrangement on products with embedded AI (ECE/TRADE/486, 2024), for which the author served as project lead and which functions as the regulatory-side counterpart to EW-AiRM™ at the enterprise level.** This lineage matters for the Committee's purposes: the framework has been stress-tested through inter-governmental negotiation and regulated-industry practice, and it is designed to interoperate with, rather than compete against, the standards and statutes that national authorities adopt.

Figure 2. The EW-AiRM™ lineage: a decade of published, intergovernmentally tested governance instruments.

SECTION 2

Executive Summary

In one paragraph. The questions before the Committee are risk governance questions. Australia currently regulates the physical inputs to data centres through planning, energy, water and environmental approval regimes, and guides the AI that runs inside them through voluntary instruments, most recently the National AI Centre's Guidance for AI Adoption and the National AI Plan's reliance on existing technology-neutral law. **What is missing is the connective layer in between: an enforceable, proportionate requirement** that the organisations building and operating AI infrastructure, and the counterparties to Government agreements with global AI companies, manage AI risk on an enterprise-wide basis, with a named accountable owner, structured pre-deployment assessment, quantified thresholds for water and energy, tested override and contingency capability, documented acceptance of residual risk, and protected incident reporting. **EW-AiRM™ (Enterprise-Wide AI Risk Management) supplies exactly this layer, and its Five Non-Negotiables, its tiered proportionality model and its Key X Indicator measurement discipline can be adapted directly into approval conditions, agreement clauses and the mandatory requirements for large AI data centres foreshadowed in the Government's July 2026 announcement.**

This submission makes three observations and eight recommendations.

- **Observation 1 (Term of Reference a).** The existing regulatory landscape is fragmented and cannot currently be judged effective, because effectiveness presupposes measurement. Facility-level water and energy data are not systematically collected or published, connection-request pipelines are distorted by speculative "phantom demand", and the April 2026 memoranda of understanding between the Government and global AI companies were concluded without a published enterprise-wide AI risk baseline against which their counterparties can be held to account.
- **Observation 2 (Term of Reference b).** The impacts on communities, industries, the environment, water and energy are identifiable and categorisable in advance. The MIT AI Risk Repository, which anchors the EW-AiRM™ operational layer, already categorises these harms, including subdomain 6.6, Environmental harm. The governance failure mode is not ignorance of the risks but the absence of thresholds, owners and monitoring obligations attached to them.
- **Observation 3 (Term of Reference c).** Concentration is the systemic risk hiding inside this Inquiry. Dependence of Australian industry and government on a small number of foreign foundation model providers and hyperscale operators corresponds to two of the eight AI Black Swan categories in the EW-AiRM™ resilience layer: systemic foundation model concentration and supply chain compromises. Resilience obligations, not only approval conditions, should follow.

"A risk appetite statement without specific thresholds is a press release, not a governance instrument."

EW-AIRM™ GOVERNANCE MAXIM · KREBSZ, ENTERPRISE-WIDE AI RISK MANAGEMENT (WILEY FINANCE, 2026)

Summary of recommendations

R1 **Enterprise-wide AI risk management as an approval condition.** Major data centre approvals and renewals should require the operator to evidence an enterprise-wide AI risk management capability with a named accountable owner for each facility and each material AI system, from deployment through decommissioning.

R2 **Non-negotiable conditions in Government agreements with global AI companies.** Existing and future agreements, including the April 2026 memoranda of understanding, should incorporate auditable minimum requirements modelled on the EW-AiRM™ Five Non-Negotiables, with residual risk formally documented and accepted at a named senior level on the Australian side.

R3 **A structured necessity and pre-deployment assessment.** Before approval, proponents should answer, in a published assessment, the three questions at the core of the HAiPECR™ pre-deployment discipline: should this be built here at this scale, how will it change behaviour and resource flows in the host community and grid, and can its impacts be explained and evidenced.

R4 **Mandatory facility-level measurement and public reporting.** Water withdrawal and consumption, energy use, water and power usage effectiveness and emissions should be measured and published per facility through a Key X Indicator regime with defined thresholds and escalation triggers, integrated with national water, energy and environmental accounts.

R5 **Concentration treated as a systemic national risk.** The Australian AI Safety Institute and system regulators should horizon-scan and stress-test dependence on individual foundation model providers, cloud regions and hyperscale operators, with documented contingency and exit arrangements for critical public-sector workloads.

R6 **Proportionate tiering.** Obligations should scale with facility size and criticality, mirroring the EW-AiRM™ Core, Standard and Full tiers, so that the forthcoming mandatory requirements for large AI data centres bind the largest operators tightly without suffocating smaller Australian providers.

R7 **Whole-of-life obligations.** Approval conditions should extend through the full lifecycle, including repurposing and decommissioning, consistent with the EW-AiRM™ strategic pillar of Through-the-Lifecycle Monitoring and Adaptability.

R8 **Protected incident reporting.** A reporting pathway for AI and data centre incidents, with retaliation protection for those who report, should be established, consistent with the fifth EW-AiRM™ Non-Negotiable and with international whistleblower protection practice.

SECTION 3 · TERM OF REFERENCE (A)

The Effectiveness of Existing Regulatory Frameworks

3.1 A landscape of parts without a whole

Australia's present arrangements are best understood as a set of vertically effective but horizontally disconnected regimes. Planning and environmental approval law governs siting; energy market rules, including the Australian Energy Market Commission's March 2026 draft connection reforms, govern grid behaviour; water utilities govern supply; and AI itself is addressed through voluntary instruments, principally the National AI Centre's Guidance for AI Adoption of October 2025 with its six essential practices, following the National AI Plan's December 2025 decision to rely on existing technology-neutral laws, sector regulators and the newly established Australian AI Safety Institute rather than economy-wide mandatory guardrails.[1,2] The Government's announcement of 15 July 2026 that it intends to legislate Australian Standards for AI, establish an Office of AI within the Department of the Prime Minister and Cabinet, and introduce mandatory requirements for large AI data centres signals welcome movement.[3]

From an enterprise-wide AI risk management perspective, the defect in the current settings is structural rather than incidental. Each regime sees one input. None requires the operator, as an organisation, to understand and govern the combined risk it carries: the AI workloads, the water, the energy, the community footprint, the supply chain and the failure modes, as one connected exposure with one accountable owner. In EW-AiRM™ terms, Australia has partial operational-layer coverage, thin strategic-layer coverage and almost no resilience-layer coverage. Governance that looks right on an organisation chart of regulators is not the same as governance that works when a hyperscale facility, its water allocation, its grid connection and its AI workloads interact in ways no single regulator observes.

3.2 Effectiveness presupposes measurement

The Committee is asked whether existing frameworks are effective. That question cannot presently be answered in the affirmative or the negative, because the measurement foundation is missing. Facility-level water and energy consumption is not systematically identified in official statistics; sector estimates suggest individual sites may seek between 5 and 40 million litres of water per day, yet the derivation of such figures is unclear.[4] Meanwhile the connection-request pipeline is distorted by speculative duplication, with industry analysis suggesting that six of every seven megawatts of requested connections are not expected to materialise.[5] A regulator cannot set thresholds against data it does not have, and a Parliament cannot evaluate effectiveness against outcomes no one measures. EW-AiRM™ addresses this through its Key X Indicator (KXI) discipline, an umbrella for key performance, key risk and key control indicators, each with defined thresholds, a named owner and an escalation path. Recommendation R4 applies that discipline at national scale.

3.3 Government agreements with global AI companies

The Terms of Reference direct particular attention to existing and future deals between the Government and global AI companies, such as the memoranda of understanding concluded with Anthropic and Microsoft in April 2026.[6] Such agreements are strategically understandable. They are also precisely the point at which enterprise-wide AI risk requirements are cheapest to introduce and hardest to retrofit. A memorandum that secures investment without securing governance transfers risk to Australian communities, ratepayers and water catchments while accountability remains offshore.

EW-AiRM™ distils enterprise AI governance to Five Non-Negotiables, reproduced in Annex A. Adapted to the intergovernmental setting, they translate into agreement clauses that any counterparty operating mature governance should already satisfy: a named accountable owner on both sides for each facility and material AI system (NN1); a structured pre-deployment assessment before capacity is committed (NN2); tested contingency and override arrangements, including grid-disturbance ride-through and workload curtailment capability, operable within defined timeframes (NN3); explicit, documented acceptance of residual risk by a named senior

Australian official rather than diffuse committee ownership (NN4); and a protected incident reporting pathway (NN5). None of these requirements is exotic. Their absence from a public agreement is itself a risk signal.

> *"Accountability cannot live in a committee: A named human being must own every deployed AI system."*
>
> EW-AIRM™ GOVERNANCE MAXIM · KREBSZ, ENTERPRISE-WIDE AI RISK MANAGEMENT (WILEY FINANCE, 2026)

3.4 Hardening voluntary guidance without a full AI Act

The Committee need not choose between the status quo and a comprehensive AI statute. The EW-AiRM™ experience in regulated industries is that a small set of auditable minimums, applied proportionately, delivers most of the assurance at a fraction of the legislative weight. The framework's tiering model, Core for small deployments, Standard for mid-sized organisations and Full for large enterprises operating many systems, offers a ready template for the proportionality the Government's foreshadowed mandatory requirements for large AI data centres will need: bind the hyperscale tier tightly, hold the middle tier to structured self-assessment with audit rights, and keep the entry tier light enough that Australian sovereign providers are not competitively disadvantaged against global incumbents. This is the same interoperability logic that underpins the UNECE Common Regulatory Arrangement (ECE/TRADE/486), which allows UN member states to converge on shared minimum requirements for products /services with embedded AI without harmonising their entire legal and regulatory systems.

SECTION 4 · TERM OF REFERENCE (B)

Impacts on Communities, Industries, the Environment, Water and Energy

The potential impacts the Committee has been asked to examine are neither speculative nor unclassifiable. The MIT AI Risk Repository, which anchors the EW-AiRM™ operational layer, categorises more than 1,700 documented AI risks across seven domains, and pairs them with 831 evidence-based controls in the associated MIT AI Risk Mitigation Database.[7] **The impacts within this Term of Reference cluster in Domain 6, Socioeconomic and Environmental harms, including subdomain 6.6, Environmental harm.** The policy significance is this: because the risks are already categorised and controls already catalogued, the barrier to governing them is not knowledge but the assignment of thresholds, owners and monitoring obligations. What follows applies the EW-AiRM™ lens to each impact category the Committee has named.

4.1 Communities

Community impact is where the HAiPECR™ inclusivity and participation dimension does its work. Inclusivity, in the EW-AiRM™ formulation, is not a separate consultation step bolted onto an approval; it is a property of every other dimension, from siting through operation to decommissioning. Applied here, that means host communities should participate in the necessity assessment itself (should this facility be built here, at this scale), not merely comment on a decision already taken; that grievance and incident pathways should be open to residents, not only employees; and that the behavioural question at the heart of the HAiPECR™ pre-deployment assessment, how does this deployment change behaviour, should expressly cover changes to local water security, electricity prices, land use, noise and employment patterns. Where communities carry the residual risk, the documented acceptance of that risk (NN4) should name the official who accepted it on their behalf.

4.2 Industries

For Australian industry, the principal exposures are dependency and displacement. EW-AiRM™ frames every organisational AI decision against two non-negotiable strategic objectives: organisational survival, and long-term sustainable benefit. Scaled to the national level, the parallel is direct. An industrial base that becomes structurally dependent on a small number of foreign foundation model providers and hyperscale operators has optimised for short-term capability at the price of survival-level concentration risk. This is not an argument against foreign investment; it is an argument for governing it with eyes open, including contingency arrangements, data portability and exit rights in public-sector and critical-infrastructure contracts, and support for sovereign capability where dependency would otherwise become irreversible. The resilience analysis in Section 5 develops this point.

4.3 Environment, water and energy

Three EW-AiRM™ instruments bear directly on the environmental questions before the Committee.

- **Rights-based framing (HAiPECR dimension R).** The framework treats environmental sustainability as a matter of rights owed to present and future generations, aligned with the UNESCO Recommendation on the Ethics of AI, principle 8, to which Australia is a signatory. Under a rights framing, water drawn from a stressed catchment to cool speculative compute capacity is not an externality to be offset; it is a claim against future generations that requires justification through the necessity assessment.
- **Quantified thresholds (KXIs).** Every material environmental exposure should carry indicators with thresholds and owners: water withdrawal and consumption per facility, water usage effectiveness, power usage effectiveness, curtailment responsiveness, and Scope 1 to 3 emissions consistent with international sustainability reporting practice. Thresholds convert policy intent into governance; without them, commitments to "sustainable data centres" remain, in the language of the EW-AiRM™ maxims, press releases rather than governance instruments.

- **Through-the-lifecycle monitoring (Pillar 6).** Environmental obligations must survive first operation. Hardware refresh cycles, workload migration, ownership changes, repurposing and eventual decommissioning all change a facility's environmental profile. Approval conditions written only for the commissioning state will be obsolete within one refresh cycle of an industry whose risk environment changes in months, not decades.

The Committee will hear the counterargument that binding environmental requirements will deter investment. A decade of regulated-industry evidence points the other way: capital is deterred by unpredictability, not by clarity. Operators with mature enterprise-wide AI risk management already measure most of what Recommendation R4 would require them to publish.

> *"Risk management, done well, does not slow innovation. It makes innovation sustainable."*
>
> EW-AIRM™ GOVERNANCE MAXIM · KREBSZ, ENTERPRISE-WIDE AI RISK MANAGEMENT (WILEY FINANCE, 2026)

SECTION 5 · TERM OF REFERENCE (C)

Other Related Matters: Systemic Concentration and National Resilience

5.1 The risk hiding inside the Inquiry

The most consequential related matter is one no individual approval decision will surface : **systemic concentration**. The EW-AiRM™ resilience layer identifies eight categories of AI Black Swan, low-probability, high-impact events that conventional risk registers systematically miss. Two apply squarely here. Systemic foundation model concentration describes an economy in which government services, financial institutions, healthcare and industry all depend on the same small set of upstream models and cloud regions, so that a single upstream failure, withdrawal, repricing or compromise propagates simultaneously across sectors that believed themselves diversified. Supply chain compromises describe the deliberate or accidental corruption of the hardware, model weights, data or dependencies flowing through that concentrated channel. Government agreements that deepen reliance on individual global providers, whatever their other merits, mechanically increase both exposures.

The governance answer is not to predict which failure will occur. It is to build the national equivalent of the three EW-AiRM™ resilience capabilities: robust foundations (diversity of suppliers, sovereign fallback capacity, contractual exit and portability rights), continuous sensing (horizon scanning and concentration metrics maintained by the Australian AI Safety Institute and system regulators), and adaptive response (tested contingency plans for the loss or degradation of a major provider, exercised rather than filed).

> *"The response to unknown unknowns is not anticipation, it is resilience."*
>
> EW-AIRM™ GOVERNANCE MAXIM · KREBSZ, ENTERPRISE-WIDE AI RISK MANAGEMENT (WILEY FINANCE, 2026)

5.2 International interoperability

Australia's choices will land in a world already converging on shared AI governance grammar: ISO/IEC 42001 management systems, the NIST AI Risk Management Framework, the EU AI Act, the UNESCO Recommendation and the UNECE Common Regulatory Arrangement. The Committee can secure a quiet but durable advantage by ensuring that whatever obligations emerge from this Inquiry are expressed in terms interoperable with those instruments, so that compliance evidence produced for Australian regulators is portable to export markets and vice versa. EW-AiRM™ was designed as exactly such a connective framework, and the author would be pleased to assist the Committee or its Secretariat in mapping proposed obligations against these international reference points.

5.3 Offer of further assistance

The author is available to give evidence via videoconference, to provide supplementary written material on any point above, and to make available the EW-AiRM™ reference tools published at www.ewairm.com, including the pre-deployment assessment instruments and the risk-control mapping against the MIT AI Risk Repository, for the Committee's use without charge.

ANNEX A

EW-AiRM™ at a Glance

A.1 The Five Non-Negotiables

The Five Non-Negotiables are the auditable minimum of enterprise AI governance under EW-AiRM™. They are deliberately few, deliberately verifiable, and designed to be adopted as contract clauses, approval conditions or supervisory expectations without further translation.

Ref	Non-Negotiable	Application to AI infrastructure and Government agreements
NN1	**Named accountable owner**	A named individual owns every deployed AI system and, by extension, every approved facility, from deployment through decommissioning. Accountability cannot live in a committee.
NN2	**Pre-deployment assessment**	Every new deployment passes a structured HAiPECR™ assessment before go-live. At its core sit three questions: should we do it; how does it change behaviour; can we explain it? For infrastructure: should it be built here at this scale; how does it change community, water and grid behaviour; can its impacts be evidenced?
NN3	**Tested human override**	A tested, documented capability for humans to override or curtail the system within defined timeframes, typically four hours or less. For facilities: demonstrated workload curtailment and grid-disturbance response.
NN4	**Documented residual risk acceptance**	The risk that remains after controls is explicitly documented and accepted by a named senior individual. In Government agreements, that individual should sit on the Australian side.
NN5	**Protected incident reporting**	A reporting pathway for AI incidents with retaliation protection for reporters, consistent with international whistleblower protection frameworks.

A.2 Framework reference

Strategic layer: six pillars

P1 Strategic Alignment and Necessity Assessment · P2 Organisational Readiness · P3 Data and Technological Maturity · P4 Risk Tolerance and Prioritisation · P5 Governance and Accountability · P6 Through-the-Lifecycle Monitoring and Adaptability.

Resilience layer: eight AI Black Swan categories

1 Emergent capabilities · 2 Systemic foundation model concentration · 3 Alignment failures · 4 Multi-modal compound effects · 5 Adversarial breakthroughs · 6 Supply chain compromises · 7 Multi-agent emergence · 8 Quantum cryptographic transition.

Operational layer

Systematic exposure mapping against the MIT AI Risk Repository (MIT FutureTech; more than 1,700 categorised risks across 7 domains and 24 subdomains, CC BY 4.0), paired with 831 evidence-based controls from the MIT AI Risk Mitigation Database.

HAiPECR ethical overlay

Seven dimensions operating across all three layers: Human oversight and autonomy · Accountability · inclusivity and participation · Protection of data, privacy and neural integrity · Ethics · Conduct risk · Rights and non-harm including intergenerational sustainability. Mapped to the ten core principles of the UNESCO Recommendation on the Ethics of AI (2021). Listed on the OECD AI Policy Observatory since April 2023.

Proportionality: three tiers

Core (fewer than 100 employees, 1 to 5 AI systems) · Standard (100 to 1,000 employees, 5 to 25 systems) · Full (more than 1,000 employees, more than 25 systems). Implementation follows a five-phase roadmap: Assess, Enhance, Build, Implement, Sustain.

Measurement: KXIs

Key X Indicators, the umbrella for key performance, key risk and key control indicators. Every material exposure carries indicators with defined thresholds, a named owner and an escalation path.

A.3 Where to learn more

The complete framework is published as Krebsz, M., *Enterprise-Wide AI Risk Management* (Wiley Finance, November 2026). Framework tools, reference material and the EW-AiRM™ knowledge resources are available at www.ewairm.com. Institutional enquiries: Human-Ai.Institute, www.human-ai.institute, contact@human-ai.institute.

References

1. Department of Industry, Science and Resources and National AI Centre, *Guidance for AI Adoption*, October 2025 (six essential practices, superseding the Voluntary AI Safety Standard, 2024).
2. Australian Government, *National AI Plan*, 2 December 2025 (reliance on existing technology-neutral law and sector regulators; establishment of the Australian AI Safety Institute).
3. Prime Minister of Australia and PM&C Office of AI, announcement of 15 July 2026 (plans to legislate Australian Standards for AI, including mandatory requirements for large AI data centres).
4. Water Services Association of Australia estimates as cited in submissions to this Inquiry (individual sites seeking 5 to 40 million litres per day).
5. Oxford Economics analysis for Amazon Web Services on speculative connection requests in the Australian data centre market (approximately six of seven requested megawatts not expected to materialise).
6. Department of Industry, Science and Resources, memoranda of understanding with Anthropic and Microsoft, April 2026.
7. Slattery et al., *The AI Risk Repository* (2024) and Saeri et al., *The AI Risk Mitigation Database* (2025), MIT FutureTech, airisk.mit.edu, licensed CC BY 4.0.

Scope note. This submission draws on publicly available material and the author's published framework. It does not constitute legal advice. The author consents to publication of this submission in full.