

The MESA-to-E-23 Crosswalk

Where the evidence for each principle has to live.

Scored against the MESA Framework™ (Maturity, Evidence, Substrate, Alignment), specification v1.1, doi:10.5281/zenodo.22109836, listed in the Defensible AI Framework Registry v1.0, doi:10.5281/zenodo.22170112. © Dr. Nabeel A. Khan.

E-23 has twelve principles under three outcomes, plus an inventory appendix. MESA scores institutional governance at four altitudes: the Regulatory Floor, the Strategic Compass, the Operational Machinery and the Technical Substrate. This crosswalk places each principle at the altitude where its evidence must exist, names the framework note that specifies that evidence, and gives one gap question per row. The pattern to notice: the principles that are satisfied by documents sit high; the principles that are examined against records sit low; and the Technical Substrate constrains every operational control that depends on runtime enforcement.

The note column marks each reference **Deposited**, with its concept DOI, or **Playbook**, meaning the authoritative treatment is a chapter of the Enterprise Playbook and no specification is deposited. Those markers are read from the framework register when this document is built; they are not typed.

E-23 PRINCIPLE	WHAT IT EXPECTS	MESA ALTITUDE	FRAMEWORK NOTE	GAP QUESTION
1.1 Organizational enablement	Reporting structures, skilled MRM staff for novel technologies, board reporting	2 Strategic Compass	CADRE (REG-07) PLAYBOOK	Who is named for AI model risk, with what budget, and when did the board last see an AI model risk report?
1.2 MRM framework	Aligned to risk appetite, defines inventory, ratings, lifecycle, external models, periodic review	2 Strategic Compass	CADRE (REG-07) PLAYBOOK MESA MRM Framework (REG-02) PLAYBOOK	Does the framework state an AI-specific risk appetite, or inherit a generic one?
1.3 Use of models	Models deployed only where they contribute reliably; retired when no longer fit	3 Operational Machinery	Five-Gate Deployment Model (REG-06), Gate 3 Approval DEPOSITED 10.5281/zenodo.22170122	Does every AI approval state the use and carry an expiry?
2.1 Model identification	Periodic survey, triage, provisional rating, evergreen inventory including decommissioned and vendor models	3 and 4	Five-Gate Deployment Model (REG-06), Gate 2 inventory entry DEPOSITED 10.5281/zenodo.22170122	Is the inventory produced from systems or maintained by hand, and are LLM and retrieval systems in it?
2.2 Model risk rating	Quantitative and qualitative criteria including autonomy, data reliability, customer impact; external models standalone	3 Operational Machinery	MESA MRM Framework (REG-02) PLAYBOOK	Does the rating method score autonomy and explainability, or only portfolio size?
2.3 Risk management intensity	Rating drives review frequency, documentation, approval authority, monitoring, limits, mitigants	3 Operational Machinery	Five-Gate Deployment Model (REG-06) DEPOSITED 10.5281/zenodo.22170122 MESA MRM Framework (REG-02) PLAYBOOK	Can you show that a higher rating changed who approved and how often it is reviewed?
3.1 Policies, procedures, controls	Robust, flexible, explicit accountability, early stakeholder involvement, independence, documentation	3 Operational Machinery	Five-Gate Deployment Model (REG-06), as lifecycle spine DEPOSITED 10.5281/zenodo.22170122	Can each lifecycle stage name its accountable person and the record it writes?
3.2 Model data	Accurate, relevant, compliant, traceable, timely; synthetic and proxy data controlled; integrated with enterprise data governance	4 Technical Substrate	AI Data Governance Framework (REG-03) PLAYBOOK Five-Gate Deployment Model (REG-06), Gate 1 DEPOSITED 10.5281/zenodo.22170122	Was lineage recorded at training time, or reconstructed afterwards?

E-23 PRINCIPLE	WHAT IT EXPECTS	MESA ALTITUDE	FRAMEWORK NOTE	GAP QUESTION
3.3 Model development	Documentation standards, expert judgment, explainability requirements, selection and monitoring criteria	3 and 4	Five-Gate Deployment Model (REG-06), Gates 1 and 2 DEPOSITED 10.5281/zenodo.22170122 Pattern Language for Production LLM Platforms, evidence by design DEPOSITED 10.5281/zenodo.22109864	Is documentation produced as evidence by the pipeline, or written as narrative after the fact?
3.4 Model review	Independent of development, commensurate with rating, triggered by events, covers AI methods, explainability, third-party components	3 Operational Machinery	Five-Gate Deployment Model (REG-06), Gate 2 Validation DEPOSITED 10.5281/zenodo.22170122	Is the validator named, independent, and able to reject?
Model approval	Suitability for intended use; affirm rating and residual risk; compensating mitigants where limitations exist. Not a numbered principle: OSFI sets model approval out as a subsection between 3.4 and 3.5.	3 Operational Machinery	Five-Gate Deployment Model (REG-06), Gate 3 DEPOSITED 10.5281/zenodo.22170122	Does the approval record bind the policy version in force at the moment of approval?
3.5 Model deployment	Quality and change control, production tests, data consistency, documented procedures, cyber and operational risk assessment, explainability communicated	4 Technical Substrate	Five-Gate Deployment Model (REG-06), Gate 4 DEPOSITED 10.5281/zenodo.22170122 Pattern Language for Production LLM Platforms, policy-as-code and trust tiers DEPOSITED 10.5281/zenodo.22109864	Has rollback been tested, and does deployment fail closed on a policy violation?
3.6 Monitoring and decommission	Standards by rating and type, thresholds, contingency, escalation, AI-specific drift and re-parametrization, disciplined retirement	3 and 4	Five-Gate Deployment Model (REG-06), Gate 5 Operation DEPOSITED 10.5281/zenodo.22170122 AI Incident Response Protocol (REG-05) DEPOSITED 10.5281/zenodo.22285057	Are the triggers armed, and can an incident be reconstructed from records alone?
Appendix 1 inventory fields	Seventeen fields per model carrying non-negligible risk	4 Technical Substrate	Defensible AI Framework Registry, evidence layer DEPOSITED 10.5281/zenodo.22170112	Which fields are produced by a system, and which are typed by a person?
B-10 third-party overlap	External models and data in scope; standalone rating	3 Operational Machinery	AI Vendor Risk Framework (REG-04) DEPOSITED 10.5281/zenodo.22170146	Has the vendor answered a model-specific questionnaire, or a security questionnaire?

Notes

Three rows carry two altitudes because the expectation is operational while the evidence is technical. That is where E-23 will be examined hardest.

Rows 2.1, 3.2, 3.5 and Appendix 1 are the substrate rows. If the answer to the gap question is that a person types it, the control exists on paper and not in the system.

The Five-Gate Deployment Model supplies the operational spine for Outcome 3. Gate 1 corresponds to design and data, Gate 2 to review, Gate 3 to approval, Gate 4 to deployment, and Gate 5 to monitoring and decommission. Each gate's passage record supplies Appendix 1 fields as a by-product rather than as a task.

Five of the notes cited in the table are deposited on Zenodo and carry concept DOIs: Five-Gate Deployment Model, Pattern Language for Production LLM Platforms, AI Incident Response Protocol, Defensible AI Framework Registry and AI Vendor Risk Framework. Three are not: REG-02, REG-03 and REG-07. Their authoritative treatment is a chapter of the Enterprise Playbook and the note column says so. A further registry entry is also undeposited and is not cited in this table: REG-09 Cross-Border AI Architecture Patterns. No deposit date is stated for any of them. The MESA Framework itself is deposited and is cited in the attribution line above rather than in a row, because it is the instrument the table scores against. AI Incident Response Protocol (REG-05) is marked Deposited on the strength of the published record. The Defensible AI Framework Registry is still at version 1.0 and still records that entry as having no deposited specification, because moving an entry is a substantive change that waits for the next registry version.

On the REG-07 name. This table names REG-07 **CADRE** (Charter, Authority, Decision rights, Records, Escalation), the canonical name in the written specification and the name carried by the live page at nabeelkhan.com/frameworks/cadre. The Defensible AI Framework Registry is still at version 1.0 and still records this entry as the AI Governance Operating Model. A canonical name change forces a major registry version, and that version has not been issued.

PROVENANCE

Method: the MESA Framework specification v1.1, doi:10.5281/zenodo.22109836, listed in the Defensible AI Framework Registry v1.0, doi:10.5281/zenodo.22170112. Sources: OSFI, Guideline E-23 Model Risk Management (2027), published 11 September 2025, effective 1 May 2027, and OSFI's Backgrounder and covering Letter. Every framework note cited is a self-deposited technical note. Zenodo assigns a DOI to what it is given and does not referee it; none of these notes has been peer reviewed. This crosswalk is an analytical mapping. It is not legal advice, it is not OSFI guidance, and no supervisor has reviewed or endorsed it. AI-use disclosure: this document was produced by Dr. Nabeel A. Khan using an AI-assisted delivery system under human direction and review. No client data was used to train any model.