



THE MIRANDOLA INSTITUTE

AI Humanist Institution Provisional Standard

A Standard for the Development of Human and Institutional Capability

in the Age of Artificial Intelligence

Status: Provisional

Conformance model: Institutional self-declaration


Standard owner: The Mirandola Institute

Version: 1.0

Publication date: September 1, 2026

Next review: No later than 12 months after publication

Designation: AI Humanist Institution



© 2026 Mirandola Institute

All Rights Reserved.

Table of Contents

1. Purpose 1

2. Status, Intended Use, and Nature of Conformance 2

3. Relationship to Existing AI Governance..... 5

4. Normative Language..... 7

5. Institutions and Scope of Conformance 8

6. Core Concepts..... 10

7. AI-Enabled Activities and Activity Boundaries..... 12

8. Human Constituency, Capability Distribution, and Human Evidence..... 16

9. Capability Materiality and Timing of Assessment..... 19

10. Capability Baseline, Intended Trajectory, and Acceptable State 21

11. Evidence, Attribution, and Implementation Proportionality 25

12. The Four Pillars and Eight Requirements 27

13. Initial Conformance and Transition for Existing AI Uses..... 37

14. Conformance Review, Signing Authority, and Public Declaration..... 38

15. Challenge, Response, and Public Accountability..... 42

16. Declaration Validity, Reaffirmation, Withdrawal, and Lapse..... 45

17. Non-Conformity 47

18. The AI Humanist Institutional Test 52

Normative Annex A: Capability Materiality Determination..... 54

Normative Annex B: Graduated Evidence and Trajectory Requirements 57

Informative Annex C: Relationship to Existing Frameworks 60

Informative Annex D: Intellectual and Empirical Basis 62

Informative Annex E: Cross-Reference to Normative Requirements 65

1. Purpose

Artificial intelligence is expanding the range of what people and institutions can know, create, analyze, coordinate, and accomplish. Improved performance does not by itself establish improved capability. An institution may produce better work with artificial intelligence while:

- weakening expertise;
- narrowing pathways through which people develop judgment;
- externalizing institutional knowledge;
- concentrating capability among incumbent experts;
- reducing the ability of responsible people to evaluate consequential work;
- or becoming increasingly dependent on machine capabilities it does not meaningfully understand or direct.

The AI Humanist Institution Provisional Standard v1.0 addresses this developmental problem. Its central question is:

Does the institution's use of artificial intelligence contribute to the durable development of human and institutional capability?

The Standard is organized around four pillars:

- Human Agency
- Institutional Integrity
- Knowledge Stewardship
- Governed Intelligence

These pillars serve a common objective:

The advancement of machine capability should become an engine for the advancement of human and institutional capability.

The Standard does not require institutions to preserve every human task or every existing capability.

It requires consequential changes in capability to be made visible, examined prospectively, and governed deliberately.

2. Status, Intended Use, and Nature of Conformance

2.1 Provisional status

This is a Provisional Standard. It translates the current intellectual framework of AI Humanism into institutional requirements while empirical research concerning capability conversion continues to develop.

The Mirandola Institute is responsible for maintenance of this Standard. The Institute shall commence a formal review of AI Humanist Institution Provisional Standard v1.0 no later than 12 months following its publication date and shall publish the outcome of that review, whether or not revision is made. The review shall consider evidence concerning:

- expertise formation;
- capability preservation, transformation, and deterioration;
- human-machine complementarity;
- apprenticeship and professional formation;
- institutional learning;
- knowledge accumulation;
- cognitive delegation;
- human authorship;
- capability distribution;
- measurement of capability across time.

Revision in response to stronger evidence is an intended property of the Standard. Following the initial formal review, the Standard should be reviewed at least every 24 months and sooner where material evidence, technological development, or experience applying the Standard warrants reconsideration.

2.2 Intended use

The Standard is intended for institutions seeking to make the development of human and institutional capability an explicit object of AI-enabled organizational practice.

It does not replace:

- legal compliance;
- responsible-AI governance;
- AI risk management;
- cybersecurity;
- privacy;
- technical assurance;
- knowledge-management systems;
- employment law;
- professional regulation.

2.3 What conformance establishes

Conformance establishes:

disciplined process fidelity around the institutional governance of capability development.

It does not establish that:

- every human capability is improving;
- every AI adoption produces a positive developmental outcome;
- no capability has deteriorated;
- every institutional decision concerning capability is substantively optimal;
- the institution has achieved an ideal state of human development.

A conforming institution may identify a negative capability trajectory. Conformance requires that the institution subject relevant AI use to a disciplined process of:

Identify → Commit → Observe → Respond

Identify means defining the activity, determining its capability materiality, and identifying affected capabilities and constituencies.

Commit means establishing the baseline or proxy baseline, intended trajectory, acceptable state, evidence, and review period before applicable results are known.

Observe means gathering evidence and determining the actual capability trajectory.

Respond means addressing material divergence and revising work, capability pathways, delegation, or future objectives where required.

This prospective structure is intended to reduce retrospective redefinition of success.

2.4 Self-declaration

Conformance with v1.0 is based on institutional self-declaration against published criteria.

The Mirandola Institute does not certify conformity with this version of the Standard. A declaration of conformity means that the declaring institution represents that it has satisfied the applicable requirements within the declared scope and remains publicly accountable for that representation under Sections 14-16.

3. Relationship to Existing AI Governance

The AI Humanist Institution Provisional Standard complements established approaches to AI governance.

It does not replace requirements concerning:

- safety;
- security;
- privacy;
- fairness;
- transparency;
- legal compliance;
- model risk;
- data governance;
- technical assurance;
- organizational AI management;
- human rights.

Different frameworks address different questions.

Societal and normative frameworks

The **UNESCO Recommendation on the Ethics of Artificial Intelligence** addresses human dignity, rights, societal well-being, education, agency, and broader normative conditions surrounding artificial intelligence.

Public-policy frameworks

The **OECD AI Principles** address trustworthy and human-centered AI and the responsibilities of governments and AI actors.

Legal and regulatory frameworks

Applicable law, including the **EU AI Act**, establishes binding requirements governing specified AI systems and uses.

Organizational AI management

ISO/IEC 42001 establishes requirements for organizational AI-management systems.

AI risk management

The **NIST AI Risk Management Framework** establishes a structured approach to governing, mapping, measuring, and managing AI risk.

Knowledge management

ISO 30401 establishes requirements for organizational knowledge-management systems.

Human impact and well-being

IEEE 7010 addresses effects of autonomous and intelligent systems on human well-being.

These frameworks may address competence, agency, knowledge, oversight, learning, accountability, or human development.

The distinguishing object of this Standard is narrower:

the longitudinal trajectory through which machine capability and AI-assisted performance contribute to human and institutional capability.

The Standard does not claim that other frameworks ignore capability. It makes **capability trajectory** the principal object of conformity.

4. Normative Language

Within this Standard:

shall indicates a requirement necessary for conformity.

should indicates a recommended practice.

may indicates a permitted option.

can indicates possibility or capability.

Capitalized classifications **Low**, **Significant**, and **High** refer exclusively to the capability-materiality classes defined in Section 9 and Annex A.

The ordinary adjective *material* means sufficiently consequential to affect a determination, conclusion, or required response.

A **Conformity condition** states the condition that shall be satisfied for the associated AH requirement to conform.

4.1 Substantive review

Substantive review means a review that examines current evidence against the applicable acceptable state, reaches a documented determination regarding capability trajectory, identifies any material divergence, and determines whether a response is required.

A substantive review is distinct from administrative confirmation that an inventory or record remains current.

5. Institutions and Scope of Conformance

5.1 Institution

Institution means the legal or organizational entity making the conformity declaration.

A subsidiary, university, public agency, nonprofit organization, business entity, or organizational unit with sufficient independent governance may constitute an institution where it possesses authority to satisfy the Standard within its declared scope.

A parent organization shall not make an enterprise-wide declaration covering separately controlled entities unless those entities are included in the:

- activity inventory;
- capability-materiality process;
- evidence system;
- review process;
- declaration authority

required by this Standard.

A subsidiary may make its own enterprise-wide declaration where the declaration covers that subsidiary as a whole.

5.2 Enterprise-wide declaration

An enterprise-wide declaration covers the declaring institution as a whole and may be made only where the requirements of the Standard have been applied across the institution as a whole in accordance with Section 5.4.

The **unqualified** designation:

AI Humanist Institution

shall be used only where the declaration is enterprise-wide.

5.3 Scoped declaration

A declaration may cover a defined:

- division;
- school;
- faculty;
- practice;
- program;
- operating unit;
- other bounded institutional scope.

A scoped declaration shall identify its scope whenever conformity is communicated. It shall not reasonably imply conformity by the institution as a whole.

5.4 Scope integrity

An enterprise-wide declaration shall include the organizational areas in which the institution's most significant capability effects from AI use occur. The institution shall not exclude an organizational area where doing so would materially change the character of the enterprise-wide capability assessment.

A scoped declaration shall identify:

- the scope included;
- significant organizational units excluded;
- the reasons for exclusion.

6. Core Concepts

6.1 Machine capability

What an intelligent system can do.

6.2 AI-assisted performance

What a person or institution can accomplish while using an intelligent system.

6.3 Human capability

The durable capacity of people to understand, judge, learn, create, communicate, decide, coordinate, direct, and act.

6.4 Institutional capability

The durable capacity of an institution to pursue purposes, combine knowledge, coordinate action, exercise judgment, preserve understanding, learn, adapt, and execute.

These constructs may develop together. They need not.

An increase in machine capability may increase AI-assisted performance without producing a corresponding increase in human or institutional capability.

6.5 Capability conversion

Capability conversion is the process through which machine capability or AI-assisted performance contributes to durable human or institutional capability.

6.6 Durable capability

Durable capability does not mean machine-independent capability. Durable capability is a reliable capacity embedded in people, knowledge, practices, relationships, institutions, or governable human-machine systems that persists beyond a transient machine output and contributes to future understanding or action.

Durable capability may exist where an institution becomes reliably better able to:

- undertake similar work again;
- understand the work being performed;
- improve through experience;
- direct the systems involved;
- recognize material exceptions or failures;
- combine machine capability with human judgment;
- transfer knowledge;
- adapt to new circumstances;
- extend capability into new forms of work.

The relevant question is not:

Could we perform this work without AI?

It is:

What reliable capacity now exists because AI became part of the work?

6.7 Human authorship

Human authorship is the meaningful human capacity to establish purposes, exercise consequential judgment, direct delegated action, evaluate results, revise course, and assume responsibility. Human authorship does not require human execution of every task. Human presence, formal approval, or nominal responsibility does not establish meaningful authorship where the responsible person lacks sufficient capability to understand, evaluate, or revise consequential work.

The governing principle is:

Delegate execution while preserving human authorship.

7. AI-Enabled Activities and Activity Boundaries

7.1 Assessable AI-enabled activity

An assessable AI-enabled activity is an identifiable use or coherent set of related uses of artificial intelligence that performs or materially changes a body of work and can plausibly affect an identifiable human capability, developmental pathway, or institutional capability.

Each activity shall possess sufficient coherence that its capability effects can reasonably be assessed together.

Illustrative activity units include:

- AI-assisted legal research and drafting by junior associates;
- AI-supported diagnostic judgment within a clinical service;
- predictive-maintenance analysis within a manufacturing operation;
- AI-assisted student writing within an academic program;
- AI-supported investment research within an investment team.

These examples illustrate activity granularity only. They do not establish a materiality classification or other conformity result. Detailed examples are provided separately in implementation guidance.

7.2 De minimis exclusion

Incidental AI functionality need not be inventoried as a separate assessable activity where it does not plausibly have a meaningful capability effect.

Examples may include ordinary:

- spellchecking;
- formatting;
- interface autocomplete;
- scheduling assistance;

- low-consequence administrative automation.

The institution shall not apply the de minimis exclusion where cumulative or repeated use could plausibly affect a consequential capability or developmental pathway.

The de minimis exclusion shall not be applied to a use of artificial intelligence that independently triggers AH-2 or AH-6 under Section 12. Such a use shall be inventoried and classified, and the requirements of Section 10.7 shall apply where it is classified Low.

7.3 Activity boundary

Activity boundaries shall be drawn so that each activity has a reasonably coherent capability effect on an identifiable:

- constituency;
- developmental pathway;
- role;
- institutional function;
- body of knowledge.

7.4 Aggregation

An institution shall not aggregate AI uses where doing so would conceal materially different capability effects across:

- roles;
- experience levels;
- career stages;
- functions;
- developmental pathways.

Where materially different effects exist, they shall be assessed separately.

7.5 Fragmentation

An institution shall not subdivide related AI use where doing so would materially understate the:

- scale;
- distribution;
- persistence;
- recoverability;
- developmental significance

of a common capability effect.

7.6 Cumulative effects

Where multiple activities individually classify Low but affect the same:

- capability;
- constituency;
- developmental pathway;
- institutional function,

the institution shall assess their cumulative effect.

Where the cumulative effect satisfies Significant or High criteria, the institution shall assess the activities collectively or through another method that captures the combined effect.

7.7 Activity inventory

The institution shall maintain an inventory of all assessable AI-enabled activities within the declared scope.

For each activity the inventory shall include:

- activity description;
- principal constituency or institutional function;
- assessment date;

- Low, Significant, or High classification;
- responsible organizational function.

The public declaration shall contain the aggregate inventory information required by Section 14.3.

8. Human Constituency, Capability Distribution, and Human Evidence

8.1 Relevant human constituency

The relevant human constituency includes:

people whose capability the institution has a material role in developing, or upon whose judgment the institution materially depends.

Depending on institutional context, this may include:

- employees;
- professionals;
- managers;
- trainees;
- students;
- researchers;
- members;
- public servants;
- other groups materially connected to the institution's capability system.

8.2 Institutional boundary

This is an institutional capability standard. A capability may remain valuable to an individual's employability, profession, personal development, or future work even where the institution determines it is no longer necessary within the declared scope. A Retire decision therefore establishes only an institutional determination. It does not establish that the capability lacks value to the individual, profession, or society.

8.3 Distribution

Aggregate improvement shall not establish conformity where it conceals material deterioration in an important:

- role;
- cohort;
- career stage;
- professional group;
- developmental pathway.

Significant and High activities shall be assessed at a level sufficient to reveal materially different capability effects.

8.4 Outsourcing and externalization

Where outsourcing or contracting materially changes the institution's:

- expertise;
- knowledge;
- capability pathways;
- access to consequential judgment;
- ability to understand or direct important work,

those effects shall be included in the assessment.

The Standard does not impose a general obligation on an institution to develop the capabilities of every contractor or external worker. Its concern is with the effect of outsourcing or externalization on the declaring institution's own capability system.

8.5 Level of human assessment

Capability assessment shall be conducted at the:

- role;
- cohort;
- pathway;
- function;

- institutional

level wherever that level is sufficient to answer the conformity question. Individual-level evidence shall not be collected where less granular evidence is reasonably sufficient.

8.6 Human-evidence protections

Capability assessment under this Standard is not designed as:

- employee ranking;
- disciplinary assessment;
- individual productivity surveillance;
- a substitute for ordinary performance management.

Where individual evidence is necessary, applicable requirements concerning privacy, employment, data protection, worker consultation, collective representation, and professional obligations shall govern.

Only the level of individual data reasonably necessary for the conformity determination shall be collected.

9. Capability Materiality and Timing of Assessment

9.1 Capability materiality

Capability materiality is the extent to which an assessable AI-enabled activity can meaningfully affect consequential human judgment, expertise, knowledge formation, developmental pathways, institutional learning, or future institutional capability.

Each assessable activity shall be classified:

- Low
- Significant
- High

under Annex A.

9.2 Timing of initial classification

An assessable AI-enabled activity shall receive an initial capability-materiality determination **before routine operational deployment where reasonably practicable.**

Where prospective classification is not reasonably practicable, classification shall occur promptly after commencement.

An institution shall not delay classification until its next annual reaffirmation where the activity has entered routine use.

9.3 Prospective record for new Significant or High activities

Where a new activity is classified Significant or High, the capability record required by Section 10 and Annex B shall be established before the first prospective assessment period begins.

The institution shall not use capability outcomes already observed from that assessment period to define the acceptable state applicable to the same period.

This requirement implements the prospective ordering required by the **Identify → Commit → Observe → Respond** cycle established in Section 2.3.

9.4 Materiality dimensions

The institution shall assess:

1. Developmental Significance
2. Consequence of Capability Loss
3. Recoverability
4. Scale and Distribution
5. Persistence

9.5 Upward-only discretion

The institution may classify an activity above the level produced by Annex A. It shall not classify an activity below that level.

9.6 No Significant or High activities

Where an institution uses AI in a core professional, educational, research, operational, managerial, or decision-making function yet identifies no Significant or High activities, its public declaration shall:

1. state that conclusion;
2. explain its basis;
3. disclose sufficient inventory information to make the conclusion reasonably contestable.

10. Capability Baseline, Intended Trajectory, and Acceptable State

10.1 Capability record

Every Significant and High activity shall maintain a capability record proportionate to its classification.

The record shall identify:

- relevant capability;
- affected constituency or function;
- baseline or proxy baseline;
- intended trajectory;
- acceptable state;
- evidence;
- assessment period;
- review date;
- responsible authority.

Annex B establishes graduated requirements for Significant and High activities.

10.2 Intended trajectory

Each relevant capability shall be designated:

Develop

Capability is intended to increase.

Preserve

Capability is intended to remain sufficient for future institutional purposes.

Transform

Capability is intended to change substantially because the human or institutional role is changing.

Retire

The institution has determined that maintaining the capability is no longer necessary within the declared scope.

10.3 Acceptable state

Develop, Preserve, and Transform capabilities shall have an acceptable state or decision criterion. The criterion shall be established **before results from the applicable assessment period are known**. For Significant activities, the acceptable state may be qualitative, quantitative, or mixed. For High activities, it shall be sufficiently specific to support a documented determination of whether material divergence occurred.

10.4 Revision

An acceptable state or intended trajectory may be revised when evidence or circumstances warrant.

The record shall state:

- previous criterion or trajectory;
- revised criterion or trajectory;
- date;
- reason;
- whether applicable results were already known.

Except for correction of a documented error, revised criteria or trajectories shall apply prospectively. An institution shall not redefine success for a completed assessment period merely because the original criterion was not achieved.

10.5 Material divergence

Where evidence indicates material divergence from a Develop, Preserve, or Transform trajectory, the institution shall determine and document an appropriate response proportionate to capability materiality.

Implementation of that response may extend beyond the assessment date.

Conformity requires that a required response has been determined and documented; it does not require that every response has been fully implemented by the time of declaration.

10.6 Governance of Retire decisions

For every Significant or High Retire decision the institution shall document:

- the capability;
- why it is no longer necessary within the declared scope;
- what capability, process, system, or arrangement now performs the relevant function;
- whether the capability remains necessary for AH-2;
- whether it remains necessary for AH-6;
- whether retirement affects another capability pathway the institution still requires;
- approving authority.

A High Retire decision shall be approved by an authority organizationally above, or otherwise independent from, the operating owner of the affected activity. Retirement shall be prospective. Unexamined deterioration shall not be redesignated retrospectively as intentional retirement. A subsequent Retire decision shall not remove, alter, or supersede the record of any material divergence identified before the Retire decision.

10.7 Independent AH-2 and AH-6 capability records for Low activities

Where AH-2 or AH-6 applies independently to an activity classified Low under Section 12, the institution shall maintain a limited capability record for the human-authorship or independent-judgment capability giving rise to the requirement.

The record shall identify:

- the capability being assessed;
- the relevant constituency, role, or institutional function;
- the acceptable state;
- evidence reasonably sufficient to assess that state;
- the review date;
- the responsible authority.

The acceptable state shall be established before results from the applicable review period are known.

The institution shall conduct a substantive review of the capability at an interval proportionate to the consequence of the human responsibility or reliance involved, and in no case less frequently than every 24 months. A substantive review shall also occur sooner where material change makes the prior assessment no longer reasonably representative. Where evidence indicates material deterioration relative to the acceptable state, the response requirements of Section 10.5 shall apply.

This limited record does not change the capability-materiality classification of the underlying activity and does not require the full Significant or High capability record specified in Annex B.

11. Evidence, Attribution, and Implementation Proportionality

11.1 Evidence

The institution shall maintain evidence reasonably sufficient to support:

- activity boundaries;
- materiality classifications;
- baselines;
- intended trajectories;
- acceptable states;
- trajectory determinations;
- Retire decisions;
- divergence responses;
- conformity conclusions.

11.2 Performance is not capability

Improvement in output quality, productivity, speed, scale, or efficiency shall not by itself establish increased human or institutional capability. **Separate evidence** shall support the capability conclusion.

11.3 Attribution

The Standard does not require exclusive causal attribution.

The institution shall provide:

a reasonable, evidence-based account of the contribution of AI-enabled work to observed capability change.

The account should distinguish:

- observed change;
- interpretation;

- plausible AI contribution;
- material alternative explanations;
- material uncertainty.

11.4 Implementation proportionality

The formality, administrative complexity, and specific mechanisms used to satisfy the Standard may be proportionate to:

- institutional size;
- organizational complexity;
- available governance structures;
- nature of the affected work.

Implementation proportionality shall not:

- reduce the materiality classification of an activity;
- eliminate a substantive requirement applicable to the activity;
- lower the evidentiary sufficiency required to support a conformity conclusion.

A small institution may satisfy a High requirement through simpler mechanisms than a large institution, provided those mechanisms produce evidence reasonably sufficient for the same substantive determination.

Scale the machinery, not the standard.

12. The Four Pillars and Eight Requirements

Significant and High activities shall be considered against all eight requirements. Where the institution determines that a requirement is not applicable to a Significant or High activity, it shall document:

- the requirement;
- the reason it is not applicable;
- the factual basis for the determination.

Strength under one requirement shall not compensate for failure under another applicable requirement.

AH-2 and AH-6 have independent applicability triggers. They apply where consequential human responsibility or consequential reliance on AI-generated claims, recommendations, analysis, or interpretations exists, including where the associated activity is classified Low.

Where AH-2 or AH-6 is independently triggered for a Low activity, the limited capability-record and substantive-review requirements of Section 10.7 shall apply.

The Standard therefore contains two distinct routes to capability scrutiny:

- a developmental trigger, where an activity is classified Significant or High because AI can materially affect capability; and
- a responsibility trigger, where human authorship or independent judgment remains consequential even though the activity itself is classified Low in capability materiality. A use giving rise to a responsibility trigger shall be inventoried and classified and shall not be excluded under Section 7.2.

Capability materiality and consequential human responsibility are related but distinct questions.

12.1 Human Agency

AH-1 - Capability Formation

Requirement

For Significant and High activities affecting human capability, the institution shall identify:

- consequential human capabilities affected;
- relevant constituencies;
- materially different effects across roles or developmental stages;
- intended developmental direction.

Section 10 governs the associated baseline, trajectory, acceptable-state, and evidence mechanics. Increased AI-assisted performance shall not by itself establish increased human capability.

Conformity condition

Material changes in consequential human capability shall be governed as developmental choices rather than occurring solely through unexamined automation.

Principle

AI adoption should change human capability intentionally rather than accidentally.

AH-2 - Sustained Human Authorship

Requirement

Where consequential human responsibility remains, the institution shall identify the capabilities necessary for meaningful human authorship and assess whether those capabilities remain sufficient across time.

The assessment shall consider, as applicable, the capacity to:

- understand the relevant purpose;
- identify material assumptions;
- examine important evidence;
- understand consequential alternatives;
- evaluate machine-generated recommendations or conclusions;
- recognize circumstances requiring reconsideration;
- revise direction;
- explain consequential judgment.

Point-in-time competence, training completion, approval authority, or override capability shall not by themselves establish conformity.

Conformity condition

Where consequential human responsibility remains, the capabilities necessary to exercise that responsibility shall remain in an acceptable state across time.

Material deterioration shall require a response under Section 10.5.

Principle

Human responsibility should remain supported by human capability across time.

12.2 Institutional Integrity

AH-3 - Institutional Capability Conversion

Requirement

For Significant and High activities, the institution shall distinguish among:

- machine capability;
- AI-assisted institutional performance;
- durable institutional capability.

The institution shall identify relevant institutional capabilities designated to Develop, Preserve, Transform, or Retire and assess them under Section 10.

Conformity condition

Improved AI-assisted output shall not be characterized as institutional capability development without evidence that relevant capability has become durably embedded in people, knowledge, practices, relationships, institutional arrangements, or governable human-machine systems.

Principle

Borrowed capability and developed capability are not the same thing.

AH-4 - Capability Pathways

Requirement

Where AI changes or removes work through which consequential future capability has historically developed, the institution shall determine whether that capability remains necessary. Where it remains necessary, the institution shall maintain or establish a credible pathway through which future people can develop it.

The assessment shall consider materially different effects across relevant:

- career stages;
- levels of experience;
- professional groups;
- developmental cohorts.

The institution shall not satisfy this requirement merely by preserving obsolete work. It shall not demonstrate conformity solely through the capabilities of incumbent experts where AI use materially undermines the pathway through which future required expertise is formed.

Conformity condition

Every consequential capability the institution continues to require shall have a credible mechanism for future formation.

Principle

Preserve the pathway to capability, not necessarily the task that once produced it.

12.3 Knowledge Stewardship

AH-5 - AI-Mediated Knowledge Accumulation

Requirement

For Significant and High activities, the institution shall identify AI-enabled:

- reasoning;
- evidence;
- decisions;
- methods;
- exceptions;
- failures;
- discoveries;
- lessons

that are materially relevant to future institutional capability.

The institution shall maintain mechanisms through which such experience can become durable human or institutional knowledge.

The institution shall distinguish:

- stored information;
- documentation;
- institutional knowledge.

Conformity condition

Material AI-enabled experience shall contribute to future institutional understanding where that understanding remains necessary for sustained institutional capability.

Principle

Execution should leave understanding behind.

AH-6 - Sustained Independent Judgment

Requirement

Where people rely consequentially on AI-produced claims, recommendations, analysis, or interpretations, the institution shall identify the level of independent judgment required for consequential use.

Relevant people **shall** remain capable, at a level appropriate to their responsibilities, of distinguishing among:

- assertion;
- evidence;
- inference;
- interpretation;
- uncertainty;
- assumption;
- speculation.

The institution shall assess whether sustained AI use Develops, Preserves, Transforms, or degrades that capability. Point-in-time AI literacy, training completion, or formal override authority shall not by itself establish conformity.

Conformity condition

Where independent judgment remains necessary to consequential work, the relevant capability shall remain in an acceptable state across time.

Material deterioration shall require a response under Section 10.5.

Principle

Fluency should expand inquiry, not terminate it.

12.4 Governed Intelligence

AH-7 - Delegation as a Learning System

Requirement

For Significant and High delegated execution, the institution shall identify execution evidence relevant to future human or institutional capability.

Such evidence may include:

- successes;
- failures;
- overrides;
- corrections;
- escalations;
- disagreements;
- exceptions;
- repeated intervention patterns;
- evidence of dependency;
- comparative human-machine strengths.

Where material lessons emerge, the institution shall maintain a mechanism through which those lessons can affect:

- human judgment;
- institutional knowledge;
- mandate design;
- task design;
- delegation;
- capability pathways;
- allocation of authority;
- human-machine coordination;
- subsequent execution.

Existing governance evidence may be used. Duplicate assurance infrastructure is not required.

Conformity condition

Where material lessons emerge from delegated execution, a mechanism shall exist through which those lessons can influence subsequent human or institutional capability or work design.

Principle

Execution should leave the institution better able to direct the next execution.

AH-8 - Capability Review and Rebalancing

Requirement

The institution shall periodically conduct substantive reviews of Significant and High capability trajectories.

The review shall consider, where applicable:

- actual versus intended trajectory;
- improvement and deterioration;
- emerging capability requirements;
- distribution across roles and developmental stages;
- effectiveness of capability pathways;
- institutional knowledge accumulation;
- sustained human authorship;
- sustained independent judgment;
- machine dependence;
- institutional ability to direct AI-enabled work.

Material divergence shall receive the documented response required by Section 10.5.

Conformity condition

The human-machine division of work shall remain subject to deliberate institutional revision rather than developing solely through technological or operational drift.

Principle

The human-machine relationship should develop through learning rather than drift.

13. Initial Conformance and Transition for Existing AI Uses

The absence of a prospective historical baseline shall not prevent initial conformity.

For an existing Significant or High activity, the institution may:

- establish the best reasonably available proxy baseline;
- reconstruct historical evidence where practicable;
- disclose material limitations;
- establish the future intended trajectory prospectively;
- establish the acceptable state before the first prospective assessment period under the Standard.

The reconstructed historical period shall not be represented as prospectively committed.

The first prospective cycle under the Standard shall provide the basis for subsequent reassessment. Where retrospective evidence is weak, the institution shall disclose that limitation rather than imply greater precision than the evidence supports.

14. Conformance Review, Signing Authority, and Public Declaration

14.1 Conformance review

Before declaration or reaffirmation, the institution shall review:

- organizational scope;
- activity inventory;
- activity boundaries;
- cumulative effects;
- materiality classifications;
- Significant and High capability records;
- AH-2 and AH-6 independent triggers;
- Section 10.7 limited capability records, acceptable states, and substantive reviews;
- non-applicability determinations;
- intended trajectories;
- Retire decisions;
- acceptable states;
- evidence;
- material divergences;
- determined responses;
- implementation status of responses;
- evidence limitations;
- challenge handling.

14.2 Signing separation

The authority approving the conformity declaration shall not be principally accountable for the AI deployment decisions under review.

Where institutional size makes such structural separation impracticable, the institution shall:

- disclose the limitation; and
- establish an independent review mechanism proportionate to its structure.

Possible mechanisms may include:

- board review;
- audit or risk committee review;
- academic governance;
- professional peer review;
- independent external review.

14.3 Public declaration

The public conformity statement shall disclose:

Scope and status

- whether conformity is enterprise-wide or scoped;
- scope covered;
- Standard version;
- assessment date;
- expiry date;
- approving authority.

Activity inventory

- number of Low activities;
- number of Significant activities;
- number of High activities;
- material exclusions;
- transitional reliance on proxy baselines where relevant.

Independent responsibility triggers

The declaration shall disclose:

- number of Low activities independently triggering AH-2;
- number of Low activities independently triggering AH-6;
- number of independently triggered Low capability records in which material divergence was identified during the reporting period;

- number of those divergences for which a response has been determined and documented.

An activity may trigger both AH-2 and AH-6; the disclosed counts therefore need not be mutually exclusive.

Capability trajectories

For Significant and High capabilities, the declaration shall disclose aggregate counts designated:

- Develop;
- Preserve;
- Transform;
- Retire.

Of the capabilities designated Retire, the declaration shall separately disclose the number designated Retire following identification of material divergence during the current or immediately preceding assessment period.

Divergence and response

The declaration shall disclose:

- number of Significant capabilities where material divergence was identified during the reporting period;
- number of High capabilities where material divergence was identified;
- number of divergences for which a response has been determined and documented;
- number of determined responses whose implementation remains in progress;
- number of divergences for which no required response has been determined.

A non-zero count of responses under implementation does not establish non-conformity. A divergence for which a response was required but has not been determined is inconsistent with Section 10.5.

Challenge handling

The declaration shall disclose aggregate counts of:

- challenges received during the reporting period;
- challenges receiving a substantive response;
- challenges determined non-qualifying under Section 15.4;
- qualifying challenges still within the 60-day response period;
- qualifying challenges for which the response period has expired without response.

The declaration need not disclose the identity of challengers or confidential details of a challenge.

Evidence

The declaration shall identify:

- material evidence limitations;
- the challenge contact point required by Section 15.1.

14.4 Permitted language

A conforming institution may state:

[Institution] has self-declared conformity with the AI Humanist Institution Provisional Standard v1.0 for the scope identified in its published conformance statement.

An enterprise-wide conforming institution may use:

AI Humanist Institution

subject to the applicable designation and mark policy.

Use of the designation is subject to the process-fidelity meaning established in Section 2.3. The institution shall not imply certification, accreditation, or independent verification where none has occurred.

15. Challenge, Response, and Public Accountability

15.1 Mandatory challenge channel

Every declaring institution shall maintain a publicly accessible contact point through which substantiated challenges to its conformity declaration may be submitted.

The challenge process shall operate independently of whether the Mirandola Institute maintains a public register.

15.2 Submission and confidentiality

A challenge should identify:

- the requirement or conformity claim at issue;
- the factual basis of the challenge;
- evidence reasonably available to the challenger.

A challenge may be submitted directly or through a representative. The challenge process **shall provide a means for confidential submission**. Anonymous submission **should be** available where reasonably practicable. The institution shall not require unnecessary disclosure of a challenger's identity.

15.3 Good-faith challenger protection

An institution shall not retaliate against a member of its relevant constituency for making, supporting, or providing evidence for a good-faith challenge under this Standard.

Retaliation inconsistent with this requirement constitutes non-conformity. This provision does not restrict an institution from addressing knowingly false, threatening, abusive, or otherwise unlawful conduct through applicable processes.

15.4 Challenge admissibility

An institution may determine that a challenge does not require substantive response where it is:

- unrelated to the Standard;
- outside the declared scope;
- unsupported by any identifiable factual basis;
- duplicative of a previously addressed challenge and presents no materially new evidence;
- abusive, threatening, or submitted for an evidently improper purpose.

The institution shall maintain a record of:

- challenges received;
- admissibility decisions;
- reasons for rejection;
- substantive responses;
- response dates.

Aggregate challenge handling shall be publicly disclosed under Section 14.3.

15.5 Response obligation

A qualifying challenge shall receive a public institutional response within **60 days** of receipt.

The institution may:

- accept the challenge and revise practice;
- revise its declaration;
- withdraw its declaration;
- dispute the challenge and provide reasons;
- explain why the challenged issue does not alter conformity.

Public responses shall protect confidential, personal, legally privileged, security-sensitive, and legitimately proprietary information.

15.6 Mirandola public register

The Mirandola Institute may maintain a public declaration register. Conformity does not depend on the existence of that register.

A Mirandola register may record:

- institution;
- declaration scope;
- Standard version;
- declaration and expiry dates;
- status;
- link to institutional disclosure;
- whether a qualifying challenge is pending;
- whether an institutional response has been received;
- whether the declaration was maintained, revised, withdrawn, or lapsed.

Mirandola should not publish unreviewed third-party allegations as factual findings. The Mirandola Institute shall not operate a public declaration or challenge register until it has published a Challenge and Register Policy governing at minimum:

- submission procedures;
- confidentiality;
- handling of anonymous or represented submissions;
- moderation;
- publication of challenge status;
- correction of inaccurate information;
- handling of disputed information;
- retention of register records.

The existence of a Challenge and Register Policy does not convert Mirandola's register into a conformity-assessment or certification process.

16. Declaration Validity, Reaffirmation, Withdrawal, and Lapse

16.1 Declaration validity

A conformity declaration shall remain valid for no more than **12 months**.

16.2 Review intervals

High activities

A substantive review shall occur at least every **12 months**.

Significant activities

A substantive review shall occur at least every **24 months**, or sooner upon material change.

Low activities

Low activities shall be reconsidered during annual inventory review and upon material change.

Where AH-2 or AH-6 applies independently to a Low activity, the capability giving rise to that requirement is additionally subject to the substantive-review interval established in Section 10.7.

New activities remain subject to Section 9.2 and shall not wait for annual review.

16.3 Annual reaffirmation

Annual reaffirmation shall confirm that:

- scope remains accurate;
- the activity inventory remains current;
- newly deployed activities were timely classified;
- activity boundaries remain reasonable;

- cumulative effects have been considered;
- classifications remain reasonable;
- required substantive reviews are current;
- material divergences have received required responses;
- applicable challenge obligations have been satisfied;
- the public declaration remains accurate.

16.4 Voluntary withdrawal

An institution may withdraw its conformity declaration at any time.

Upon withdrawal:

- current use of the designation shall cease;
- the institution's conformity page shall identify the declaration as Withdrawn and state the effective date;
- any Mirandola register may record the declaration as Withdrawn;
- historical statements may remain available where clearly identified as historical.

Withdrawal does not by itself constitute a determination by Mirandola that the institution was non-conforming.

16.5 Lapse

A declaration lapses automatically at expiry unless reaffirmed. After lapse, an institution shall not represent itself as currently conforming. Historical statements may remain if clearly identified with their period of validity. Rules concerning continued use of any visual designation or mark shall be governed separately.

17. Non-Conformity

The following conditions are inconsistent with conformity. Clause references identify the upstream normative requirement.

Scope and institutional claims

- failure to include separately controlled entities represented as covered by an enterprise declaration - **Section 5.1**
- use of the unqualified AI Humanist Institution designation for a non-enterprise declaration - **Section 5.2**
- use of a scoped declaration in a manner that reasonably implies enterprise-wide conformity - **Section 5.3**
- exclusion of organizational areas where doing so materially changes the character of an enterprise-wide capability assessment - **Section 5.4**
- failure of a scoped declaration to identify significant organizational units excluded from scope and the reasons for their exclusion - **Section 5.4**

Activity identification and boundaries

- failure to maintain the required activity inventory - **Section 7.7**
- inappropriate use of the de minimis exclusion - **Section 7.2**
- application of the de minimis exclusion to a use of artificial intelligence that independently triggers AH-2 or AH-6 - **Section 7.2; Section 12**
- activity boundaries that do not permit coherent capability assessment - **Section 7.3**
- aggregation that conceals materially different capability effects - **Section 7.4**
- fragmentation that materially understates common capability effects - **Section 7.5**
- failure to assess cumulative effects of related Low activities where required - **Section 7.6**

Human constituency and evidence

- aggregate capability assessment that conceals material deterioration in an important role, cohort, career stage, professional group, or developmental pathway - **Section 8.3**

- failure to include material effects of outsourcing or externalization on the institution's own capability system - **Section 8.4**
- unnecessary collection of individual-level capability evidence where less granular evidence is reasonably sufficient - **Section 8.5**
- collection or use of individual-level evidence contrary to the protections required by the Standard or applicable obligations - **Section 8.6**

Materiality and prospective timing

- failure to classify an assessable activity - **Section 9.1**
- failure to classify a new activity within the timing required by the Standard - **Section 9.2**
- observing a prospective assessment period before establishing the required capability record and acceptable state - **Section 9.3**
- failure to assess all required materiality dimensions - **Section 9.4**
- assigning a classification below the level produced by Annex A - **Section 9.5; Annex A**
- AI use in a core institutional function accompanied by zero Significant or High classifications without the required disclosure and rationale - **Section 9.6**

Capability records, trajectory, and response

- failure to maintain the required Significant or High capability record - **Section 10.1; Annex B**
- failure to designate an intended trajectory - **Section 10.2**
- failure to establish the required acceptable state before applicable results are known - **Section 10.3**
- retrospective redefinition of a failed acceptable state or trajectory contrary to the prospective-revision rule - **Section 10.4**
- material divergence for which a required response has not been determined and documented - **Section 10.5**
- a Significant or High Retire decision without the required rationale, dependency assessment, pathway analysis, or approval - **Section 10.6**

- retrospective redesignation of unexamined deterioration as intentional retirement - **Section 10.6**
- alteration, removal, or supersession of a recorded material divergence through a subsequent Retire decision - **Section 10.6**
- failure to maintain the limited capability record required where AH-2 or AH-6 independently applies to a Low activity - **Section 10.7**
- failure to establish the acceptable state required for an independently triggered Low AH-2 or AH-6 capability before applicable review results are known - **Section 10.7**
- failure to conduct the substantive review required for an independently triggered Low AH-2 or AH-6 capability - **Section 10.7**
- material deterioration identified under an independently triggered Low AH-2 or AH-6 capability for which the response required by Section 10.5 has not been determined and documented - **Section 10.7; Section 10.5**

Evidence and attribution

- failure to maintain evidence reasonably sufficient to support a required conformity determination - **Section 11.1**
- treating improved AI-assisted performance alone as evidence of increased human or institutional capability - **Section 11.2; AH-1; AH-3**
- absence of a reasonable evidence-based account of AI's contribution where capability change is asserted - **Section 11.3**
- use of implementation proportionality to reduce a classification, eliminate a substantive requirement, or weaken evidentiary sufficiency - **Section 11.4**

Substantive requirements

- unsupported determination that an AH requirement is not applicable - **Section 12 preamble**
- strength in one requirement used to offset failure under another applicable requirement - **Section 12 preamble**
- consequential human capability changing solely through unexamined automation without the developmental governance required by AH-1 - **AH-1**

- consequential human responsibility remaining where the capabilities necessary for meaningful authorship are no longer in an acceptable state and no required response has been determined - **AH-2; Section 10.5**
- institutional capability development claimed solely from assisted output without evidence of durable embedding - **AH-3**
- a consequential capability the institution continues to require lacking a credible future formation pathway - **AH-4**
- preservation of obsolete work solely to create apparent conformity with AH-4 - **AH-4**
- reliance solely on incumbent expertise where AI materially undermines the pathway through which future required expertise is formed - **AH-4**
- failure to convert material AI-enabled experience into future institutional understanding where that understanding remains necessary - **AH-5**
- consequential reliance on AI-generated claims where relevant people no longer possess the independent judgment required by AH-6 - **AH-6**
- required independent-judgment capability falling outside its acceptable state without the response required by Section 10.5 - **AH-6; Section 10.5**
- material lessons emerging from delegated execution without a mechanism through which those lessons can influence subsequent capability or work design - **AH-7**
- failure to conduct the substantive capability reviews required by AH-8 and Section 16.2 - **AH-8; Section 16.2**
- allowing the human-machine division of work to develop solely through technological or operational drift without deliberate institutional review - **AH-8**

Transition, review, and declaration

- representation of a reconstructed historical period as though it had been prospectively committed - **Section 13**
- material retrospective-evidence limitations not disclosed during initial transition - **Section 13**
- failure to complete the required conformance review before declaration or reaffirmation - **Section 14.1**
- failure to maintain required signing separation or the required independent-review fallback - **Section 14.2**

- material omission from the disclosures required in the public conformity statement - **Section 14.3**
- implication of certification, accreditation, or independent verification where none occurred - **Section 14.4**

Challenges and accountability

- failure to maintain the required public challenge contact point - **Section 15.1**
- failure to provide a confidential challenge-submission mechanism - **Section 15.2**
- requiring unnecessary disclosure of a challenger's identity - **Section 15.2**
- retaliation against a good-faith challenger within the relevant constituency - **Section 15.3**
- failure to maintain the challenge-handling record required by Section 15.4 - **Section 15.4**
- failure to disclose aggregate challenge handling as required - **Section 15.4; Section 14.3**
- failure to respond publicly to a qualifying challenge within 60 days - **Section 15.5**

Reaffirmation, withdrawal, and lapse

- issuance or maintenance of a conformity declaration with a validity period exceeding 12 months - **Section 16.1**
- failure to conduct substantive reviews within the required intervals - **Section 16.2**
- reaffirmation without confirming the matters required by Section 16.3 - **Section 16.3**
- failure, following voluntary withdrawal, to identify the declaration as Withdrawn on the institution's conformity page and state the effective date - **Section 16.4**
- continued current use of the designation after voluntary withdrawal - **Section 16.4**
- continued representation of current conformity after declaration lapse - **Section 16.5**

18. The AI Humanist Institutional Test

The Standard can ultimately be expressed through four questions.

Human Agency

What are people becoming more capable of understanding, judging, creating, or doing?

Institutional Integrity

What is the institution becoming more capable of accomplishing reliably because artificial intelligence has become part of its work?

Knowledge Stewardship

What greater knowledge and understanding remain because the work was undertaken?

Governed Intelligence

What does delegated execution teach the institution about how people and machines should work together next time?

Together:

What greater durable human or institutional capability exists because machine capability became available?

An AI Humanist Institution does not preserve human work merely because it is human. It does not equate machine dependence with human decline. It does not treat productivity as evidence of development. It does not define success after seeing the outcome.

It identifies consequential capability, determines what it intends to Develop, Preserve, Transform, or Retire, establishes the basis on which that choice will be evaluated, observes the resulting trajectory, and responds when the evidence no longer supports its intended state.

The advancement of machine capability should become an engine for the advancement of human and institutional capability.

Normative Annex A: Capability Materiality Determination

A.1 Purpose

Each assessable AI-enabled activity shall be evaluated across five dimensions.

Each dimension shall be rated:

- Low;
- Moderate;
- High.

A.2 Developmental Significance

Low: The activity does not meaningfully contribute to consequential capability formation.

Moderate: The activity contributes meaningfully, but credible alternative pathways are readily available.

High: The activity is a major mechanism through which consequential capability develops.

A.3 Consequence of Capability Loss

Low: Deterioration would have limited institutional effect.

Moderate: Deterioration would impair an important function but remain reasonably manageable.

High: Deterioration could materially impair consequential judgment, institutional purpose, professional performance, resilience, or future execution.

A.4 Recoverability

Low: Capability can be reconstructed quickly and at relatively low cost.

Moderate: Reconstruction requires meaningful time, resources, training, or experience.

High: Reconstruction would require substantial accumulated experience, apprenticeship, tacit knowledge, institutional rebuilding, or extended time.

A.5 Scale and Distribution

Low: Effects are narrow and do not materially alter an important capability pathway or institutional function.

Moderate: Effects concern a meaningful population, function, or pathway.

High: Effects are widespread or materially alter a major source of future expertise or institutional capability.

A.6 Persistence

Low: Effects are temporary or readily reversible.

Moderate: Effects may persist without intervention.

High: Effects are likely to become structurally embedded or difficult to reverse.

A.7 Overall classification

Dimension pattern	Required classification
No High and 0-1 Moderate	Low
No High and 2-3 Moderate	Significant
One High, fewer than 2 additional Moderate	Significant
Four or more Moderate	High
One High plus 2 or more Moderate	High
Two or more High	High

Any single High establishes Significant capability materiality at minimum. The institution may classify upward. It shall not classify downward.

A.8 Cumulative classification

Multiple related Low activities affecting the same capability, constituency, pathway, or function shall be considered cumulatively.

The institution shall not rely upon individual Low classifications where the combined effect meets Significant or High criteria.

Normative Annex B: Graduated Evidence and Trajectory Requirements

B.1 Purpose

Significant and High activities are governed by the same substantive principles but differ in evidentiary depth.

Implementation may also be proportionate to institutional size and complexity under Section 11.4.

B.2 Significant activities

A Significant activity may use:

- a reasonable proxy baseline;
- qualitative, quantitative, or mixed acceptable-state criteria;
- role-, cohort-, pathway-, or function-level evidence;
- a simplified capability record;
- one principal evidence source where reasonably sufficient.

The record shall include at minimum:

- activity;
- capability;
- relevant constituency or function;
- baseline or proxy;
- intended trajectory;
- acceptable state;
- principal evidence;
- review date;
- responsible authority.

Substantive review shall occur at least every 24 months or upon material change.

B.3 High activities

A High activity shall receive deeper assessment.

Where reasonably practicable:

- a prospective baseline shall precede new deployment or major redesign;
- acceptable-state criteria shall support a clear divergence decision;
- materially different effects across relevant cohorts shall be examined explicitly;
- multiple or triangulated evidence sources shall be used;
- uncertainty shall be documented.

The High capability record shall include:

- activity and activity-boundary rationale;
- capability;
- constituency or institutional function;
- baseline and limitations;
- intended trajectory;
- acceptable state;
- evidence sources;
- distributional analysis;
- assessment period;
- review date;
- operating authority;
- reviewing authority;
- material uncertainty.

These record requirements are subject to the implementation-proportionality principle in Section 11.4. A smaller or less complex institution may combine fields or use simpler documentation where the resulting record remains sufficient to support each required determination.

Substantive review shall occur at least annually.

B.4 Evidence insufficient by itself

The following shall not by themselves establish capability development:

- training hours;
- licenses;
- prompt counts;
- AI usage;
- course completion;
- adoption rates;
- immediate productivity gains.

They may contribute to a broader evidence set.

B.5 Human evidence

Role-, cohort-, pathway-, function-, or institutional-level evidence shall be preferred where reasonably sufficient.

B.6 Attribution

Reasonable evidence of AI's contribution is required. Exclusive causal proof is not.

Informative Annex C: Relationship to Existing Frameworks

Framework	Principal domain	Illustrative structure	Relationship to this Standard
UNESCO Recommendation on the Ethics of AI	Human rights, dignity, societal well-being, education and human development	Normative principles	Societal and normative environment
OECD AI Principles	Trustworthy, human-centered AI	Principles for AI actors and governments	Complementary public-policy foundation
EU AI Act	Legal obligations	Article 4 - AI literacy; Article 14 - human oversight	Legal floor where applicable
ISO/IEC 42001	Organizational AI management	AI management system	Management foundation
NIST AI RMF	AI risk management	Govern, Map, Measure, Manage	Risk-management foundation
ISO 30401	Knowledge management	Knowledge-management system	Foundation relevant to AH-5
IEEE 7010	Human well-being impacts	Well-being assessment	Adjacent positive-outcome framework
AI Humanist Institution Provisional Standard v1.0	Longitudinal capability conversion	Identify → Commit → Observe → Respond	Developmental capability standard

C.1 Point-in-time competence versus capability trajectory

The EU AI Act’s Articles 4 and 14 address AI literacy and human oversight.

This Standard asks an additional question:

Is the capability underlying meaningful oversight developing, remaining viable, transforming appropriately, or deteriorating through sustained AI use?

Point-in-time competence and longitudinal capability trajectory are related but distinct objects.

C.2 Knowledge management versus AI-mediated capability conversion

Knowledge-management systems can preserve, organize, and circulate organizational knowledge.

AH-5 asks:

When cognition or execution is delegated to machines, does relevant experience become human and institutional knowledge, or remain externalized in machine performance?

Informative Annex D: Intellectual and Empirical Basis

The problems addressed by this Standard predate generative artificial intelligence.

The Standard does not claim to have discovered automation-induced capability deterioration, cognitive delegation, expertise formation, or organizational learning.

Its contribution is to make the **trajectory of human and institutional capability under sustained AI use an explicit object of institutional governance and public conformity claims.**

D.1 Automation and human capability

Bainbridge, L. (1983). *Ironies of automation*. *Automatica*, 19(6), 775-779.

Parasuraman, R., & Riley, V. (1997). *Humans and automation: Use, misuse, disuse, abuse*. *Human Factors*, 39(2), 230-253.

Parasuraman, R., & Manzey, D. H. (2010). *Complacency and bias in human use of automation: An attentional integration*. *Human Factors*, 52(3), 381-410.

D.2 Expertise and skill formation

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). *The role of deliberate practice in the acquisition of expert performance*. *Psychological Review*, 100(3), 363-406.

Cunha, F., & Heckman, J. J. (2007). *The technology of skill formation*. *American Economic Review*, 97(2), 31-47.

D.3 Organizational learning and intelligent technology

Argyris, C., & Schön, D. A. (1978). *Organizational Learning: A Theory of Action Perspective*. Addison-Wesley.

Zuboff, S. (1988). *In the Age of the Smart Machine: The Future of Work and Power*. Basic Books.

The automate/informate distinction is directly relevant to AH-5 and AH-7: technological execution can simultaneously generate information capable of enlarging organizational understanding.

D.4 Technology, tasks, and human work

Autor, D. H., Levy, F., & Murnane, R. J. (2003). *The skill content of recent technological change: An empirical exploration*. *Quarterly Journal of Economics*, 118(4), 1279-1333.

Acemoglu, D., & Restrepo, P. (2019). *Automation and new tasks: How technology displaces and reinstates labor*. *Journal of Economic Perspectives*, 33(2), 3-30.

D.5 Cognitive delegation

Risko, E. F., & Gilbert, S. J. (2016). *Cognitive offloading*. *Trends in Cognitive Sciences*, 20(9), 676-688.

Cognitive offloading describes the use of external action or tools to reduce internal cognitive demand.

This Standard does not presume that cognitive offloading is harmful. It treats offloading as one mechanism through which the distribution of cognitive work may change and therefore as a potential object of capability inquiry.

D.6 Human-AI performance

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.

Dell’Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayner, L., Candelon, F., & Lakhani, K. R. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. **Harvard Business School Working Paper No. 24-013**, subsequently published in **Organization Science**.

These studies reinforce the importance of distinguishing improved AI-assisted performance from durable capability.

D.7 Distinctive institutional proposition

The Standard's claim is not that automation can affect human skill.

Its narrower proposition is:

Capability formation, preservation, transformation, and erosion should become explicit objects of institutional governance when increasingly capable machines participate in cognition and execution.

The Standard operationalizes that proposition across:

- human agency;
- institutional capability;
- developmental pathways;
- knowledge;
- delegated execution;
- distribution;
- prospective commitment;
- longitudinal review;
- public contestability.

It therefore treats:

capability trajectory as an institutionally governed and publicly contestable object.

Informative Annex E: Cross-Reference to Normative Requirements

This Annex is informative and creates no conformity requirement. Its purpose is to show the relationship among normative obligations, conformity mechanisms, and the non-conformity conditions identified in Section 17. In the event of any inconsistency between this Annex and the normative provisions of the Standard, the normative provisions govern.

Normative object	Required action or state	Conformity path	Corresponding Section 17 failure
Scope - Section 5	Declare accurately bounded enterprise or scoped conformity	Section 14.1 review → Section 14.3 disclosure	Scope and institutional-claim failures
Activity unit - Sections 7.1-7.6	Define coherent activity units without masking through aggregation or fragmentation	Inventory → materiality → Section 14.1 review	Activity-boundary failures
De minimis and responsibility trigger - Section 7.2; Section 12	Inventory and classify uses triggering AH-2 or AH-6 rather than excluding them	Inventory → Section 10.7 record	De minimis exclusion misapplied
Inventory - Section 7.7	Maintain inventory of assessable AI-enabled activities	Section 14.1 review → Section 14.3 aggregate disclosure	Missing inventory
Distribution - Section 8.3	Reveal materially different capability effects	AH-1/AH-4 assessments → substantive review	Hidden cohort/pathway deterioration
Outsourcing - Section 8.4	Assess effects on institution's own capability system	Section 14.1 review	Omitted externalization effects
Human evidence - Sections 8.5-8.6	Use least-granular sufficient evidence and applicable protections	Evidence review → Section 14.1	Excessive/impermissible individual assessment
Classification - Section 9; Annex A	Classify activities prospectively and at required tier	Inventory → Section 14.1 → annual reaffirmation	Missing, late, or downward classification
Prospective commitment - Section 9.3; Section 10	Set trajectory and acceptable state before results	Capability record → substantive review	Retrospective target setting

Normative object	Required action or state	Conformity path	Corresponding Section 17 failure
Capability records - Section 10.1; Annex B	Maintain proportionate Significant/High record	Section 14.1 review	Missing record
Independent Low AH-2/AH-6 record - Section 10.7	Maintain limited record and acceptable state where authorship or independent judgment independently triggers on a Low activity	Limited record → substantive review → Section 10.5 response	Missing record, review, acceptable state, or required response
Trajectory - Section 10.2	Develop / Preserve / Transform / Retire	Section 14.3 aggregate disclosure	Missing trajectory
Acceptable state - Section 10.3	Pre-establish decision criterion	Substantive review	Undefined/post-hoc acceptable state
Revision - Section 10.4	Revise prospectively and document changes	Capability record	Retrospective redefinition
Divergence - Section 10.5	Determine and document response	Section 14.3 divergence/response counts	Required response absent
Retire - Section 10.6	Document rationale, dependencies, pathway effects, and approval; preserve the record of any preceding material divergence	Section 14.1 review → Section 14.3 aggregate Retire and post-divergence Retire disclosure	Ungoverned Retire decision or alteration of a preceding divergence record
Evidence - Section 11.1	Maintain evidence reasonably sufficient for determinations	Section 14.1 conformance review	Insufficient evidence
Performance/capability distinction - Section 11.2	Do not infer capability from performance alone	AH-1/AH-3 conformity tests	Performance treated as capability
Attribution - Section 11.3	Provide reasonable contribution account	Evidence record → substantive review	Unsupported AI contribution claim
Proportionality - Section 11.4	Scale implementation, not substantive burden	Annex B implementation	Requirement weakened by institution size
AH applicability - Section 12	Consider all eight; justify N/A; independently trigger AH-2/AH-6 where required	Section 14.1 review	Unsupported N/A determination

Normative object	Required action or state	Conformity path	Corresponding Section 17 failure
AH-1	Govern consequential human capability development deliberately	Section 10 trajectory record → substantive review	Unexamined human-capability change
AH-2	Sustain capability necessary for meaningful authorship	Section 10 record or Section 10.7 limited record → Acceptable state → substantive review → divergence response	Nominal responsibility without viable authorship
AH-3	Demonstrate durable institutional capability rather than assisted performance	Capability evidence → substantive review	Unsupported institutional-capability claim
AH-4	Maintain credible pathway for capabilities still required	Pathway evidence → substantive review	Required capability without future formation route
AH-5	Convert relevant AI experience into future understanding	Knowledge mechanism → review	Necessary understanding remains externalized
AH-6	Sustain independent judgment required for consequential reliance	Section 10 record or Section 10.7 limited record → Acceptable state → substantive review → response	Independent judgment no longer viable
AH-7	Use material execution lessons developmentally	Learning mechanism → substantive review	Material execution lessons cannot influence future work
AH-8	Review and rebalance human-machine work deliberately	Substantive review → response	Operational/technological drift
Transition - Section 13	Distinguish retrospective reconstruction from prospective commitment	Initial declaration disclosure	Retrospective period misrepresented
Conformance review - Section 14.1	Review entire conformity system	Signing authority	Declaration without completed review
Signing separation - Section 14.2	Separate approving authority from deployment accountability	Public disclosure / independent fallback	Self-review without required separation

Normative object	Required action or state	Conformity path	Corresponding Section 17 failure
Public declaration - Section 14.3	Publish required scope, trajectory, divergence, challenge and evidence information	Annual reaffirmation	Material disclosure omission
Challenge channel - Sections 15.1-15.2	Maintain public and confidential route	Section 14.3 contact disclosure	No usable challenge route
Challenger protection - Section 15.3	No retaliation for good-faith challenge	Challenge records / governance	Retaliation
Challenge handling - Sections 15.4-15.5	Record admissibility and respond within 60 days	Section 14.3 aggregate counts	Rejection/response process not followed
Review intervals - Section 16.2	Conduct substantive reviews on schedule	Annual reaffirmation	Overdue substantive review
Reaffirmation - Section 16.3	Reconfirm current conformity annually	New public declaration	Unsupported renewal
Withdrawal/lapse - Sections 16.4-16.5	Update public status and cease current conformity claims when a declaration is withdrawn or lapses	Public conformity page / register status where applicable	Failure to identify withdrawal or continued use after withdrawal/lapse

E.1 Scope of the conformity audit

This Annex cross-references the substantive obligations directed to institutions making conformity declarations under this Standard. Every institutional non-conformity category in Section 17 shall have an identifiable upstream normative basis. Conversely, every substantive institutional obligation expressed using *shall* terminates in an identifiable conformity mechanism, which may include:

1. a documented institutional record;
2. a substantive capability review;
3. a required institutional response;
4. a public conformity disclosure;
5. a challenge or accountability obligation;
6. a Section 17 non-conformity consequence.

Not every use of *shall* in the Standard is an institutional conformity obligation. Obligations directed to the Mirandola Institute as owner or maintainer of the Standard, including the review obligation in Section 2.1 and any register-policy obligation in Section 15.6, are governance obligations of the Standard owner and do not create institutional non-conformity under Section 17. Recommendations expressed as *should*, permissions expressed as *may*, explanatory provisions, and informative annex material likewise do not require a Section 17 consequence.

In the event that this Annex becomes inconsistent with the normative text through later amendment, the normative text governs and this Annex shall be corrected as part of the next document-maintenance revision.



The humanist project continues.



THE MIRANDOLA INSTITUTE
mirandolainstitute.org