

AI SERVICE QUALITY STANDARD

DQI_AIS1

First Edition — November 2026

DIGITAL SERVICE QUALITY INITIATIVE

Table of Contents

Foreword 4

Introduction 4

1. Scope 5

2. Normative References 5

3. Terms and Definitions 6

3.1 AI service 6

3.2 AI Service Quality Object 6

3.3 Assistance limit 6

3.4 AI model 6

3.5 AI Service Quality Object Description 7

4. AI Service Quality Object 7

4.1 Intended Functions and Results 7

4.2 Life Cycle, Essential Work, Results, and External Contributions 8

4.3 Boundaries, Interfaces, and Constituents 9

5. Quality Factors and Indicators 11

5.1 Factor 1 — Object's Functional Requirements 11

5.2 Factor 2 — Object's Technical Requirements 16

5.3 Factor 3 — Conformity Assessment 21

5.4 Factor 4 — Structural and Functional Integrity 23

5.5 Factor 5 — Harm, Risk, and Safety 28

5.6 Factor 6 — Object and Environment 33

5.7 Factor 7 — Functional Reliability 38

5.8 Factor 8 — Human and Object Interaction 41

5.9 Factor 9 — Competence of Personnel 45

5.10 Factor 10 — Regulatory and Standards Compliance 48

5.11 Factor 11 — Object and Stakeholders 51

5.12 Factor 12 — Object Monitoring 56

5.13 Factor 13 — Documentation and Traceability 59

5.14 Factor 14 — Object Adaptability 63

6. Examination Basis 64

6.1 Identification of the AI Service Quality Object 64

6.2 Applicable Basis 64

6.3 Examination Scope and Conditions 65
7. Evidence 66
8. Findings and Synthesis 67
9. Quality Claims 68
Bibliography 69
Copyright and Use Notice 70

Foreword

This Standard applies the Quality Institute's science of quality and Whole-Quality methodology to AI services. The book Whole Quality of Complex Objects explains this approach. DQI VOC1, Digital Service Quality — Fundamentals and Vocabulary, provides shared terms for describing digital services.

Safety is an essential part of Quality. An unmet requirement is a Quality Deficiency that creates a vulnerability and may lead to Harm.

This Standard is voluntary. Its application supports Evidence-based Findings and bounded Quality Claims under stated examination conditions.

Introduction

An AI service realizes its Intended Functions and Intended Results through coordinated contributions from service recipients, AI models, software, data, tools, personnel, and supporting resources. Its Work may include clarifying needs, generating information, making recommendations, supporting decisions, performing actions, and following through. Work and Results from external organizations may also be necessary for the service's Function Realization.

The Quality Object is identified by the Intended Functions and Intended Results being examined. Its Boundaries and Interfaces account for the Constituents, dependencies, and external contributions needed for their realization. Depending on the service, the object may extend across applications, systems, providers, and organizational arrangements. The recipient's participation and its relationship to the object require explicit consideration.

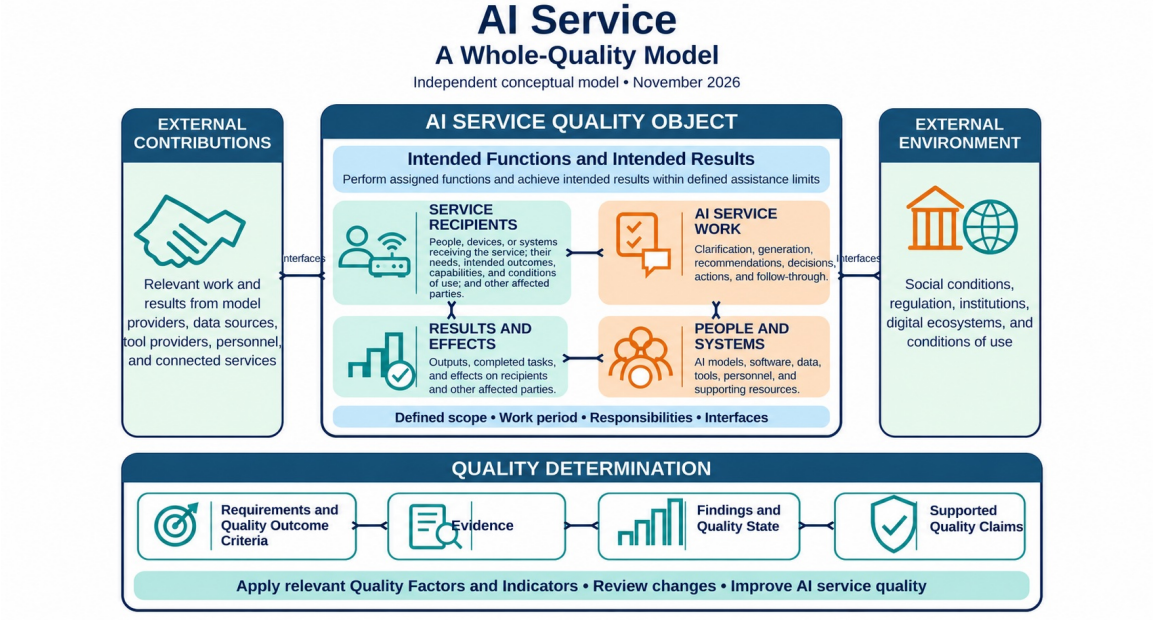


Figure 1 — AI Service Whole-Quality Model

Figure 1 presents the AI Service Quality Object, relevant external contributions and environment, and the quality-determination process. It is a conceptual model; the object's actual scope, responsibilities, and interfaces depend on the identified service.

Conditions may change during an interaction or across repeated interactions. New information, changes in recipient intent or circumstances, unavailable external contributions, and effects of earlier assistance can alter what assistance is justified. Requirements therefore address applicable outputs, sequences of interaction, completed actions, and relevant cumulative effects, including how assistance limits are established, revised, and applied.

This Standard connects the service's Intended Functions and Intended Results with applicable Object Quality Requirements and Quality Determination Requirements. Examination of the service's actual condition and performance provides Evidence for Findings and their synthesis, supporting a bounded Quality Claim.

Clause 4 establishes the AI Service Quality Object. Clause 5 states requirements under the Quality Factors and Indicators. Clauses 6 through 9 address the examination basis, Evidence, Findings and Synthesis, and Quality Claims.

1. Scope

This Standard establishes Object Quality Requirements and Quality Determination Requirements for identified AI services across relevant life-cycle stages.

It applies to AI services that provide information, generate content, support decisions, perform actions, or provide a defined combination of these functions. The identified object may involve one interaction, a sequence of interactions, or continuing service provision.

The Standard addresses the Constituents, Work and Results, Boundaries, Interfaces, dependencies, recipient participation, and external contributions necessary to realize the service's Intended Functions and Intended Results.

Requirements address functional intent, assistance limits, changing conditions, and relevant effects on recipients and other affected parties.

The Standard also establishes requirements for applying the Quality Factors and Indicators, examining Evidence, establishing and synthesizing Findings, and communicating bounded Quality Claims.

Applicable laws, regulations, professional obligations, and domain-specific requirements may contribute to the Applicable Basis.

2. Normative References

The following document is normatively referenced in this Standard and is indispensable for its application.

DQI VOC1, Digital Service Quality — Fundamentals and Vocabulary.

3. Terms and Definitions

For the purposes of this Standard, the terms and definitions given in DQI VOC1, Digital Service Quality — Fundamentals and Vocabulary, and the following apply.

This Standard specifies requirements for the AI Service Quality Object and its responsible parties, including the Evidence needed to support Quality Claims. The companion AI Service Conformity Assessment Guide addresses examination methods, sampling, assessor competence and impartiality, and the evaluation and reporting of Findings for internal and external assessments.

3.1 AI service

Digital service in which one or more AI models perform or support Work directed toward Intended Results within the functions assigned to the service.

Note 1 to entry. Work may be completed in a single interaction or continue across several interactions, chats, or periods of service provision. Individual interactions may serve separate Intended Results or produce intermediate Results that contribute to a shared Intended Result.

Note 2 to entry. The functions assigned to the service establish what Work it can undertake. The function to be performed and the result sought for a particular task may be established, clarified, or revised during service provision.

Note 3 to entry. Completion of the Work assigned to an individual interaction may precede achievement of the Intended Result of the continuing Work.

3.2 AI Service Quality Object

Identified AI service considered as an integrated whole for Quality Determination, with Boundaries established in relation to its Intended Functions and Intended Results.

3.3 Assistance limit

Boundary on the assistance an AI service may provide under specified conditions, including any prerequisites for providing that assistance.

Note 1 to entry. A limit may apply to the outputs the service provides, the decisions it supports, the actions it performs or facilitates, or the follow-through it undertakes. It may depend on demonstrated capability, available information, authorization, recipient participation, or necessary external contributions.

3.4 AI model

Computational representation that uses patterns learned from data, encoded knowledge or rules, or a combination of these to infer or generate outputs from inputs.

Note 1 to entry. Outputs may include classifications, predictions, generated content, recommendations, or proposed actions.

Note 2 to entry. One or more AI models may contribute to an Intended Function together with software, data, tools, personnel, and external contributions. The AI Service Quality Object Description identifies the model’s role and the contributions on which its use depends.

3.5 AI Service Quality Object Description

Documented representation of an identified AI service that sets out its Intended Functions and Intended Results, Boundaries, Constituents, Interfaces, and the Work and contributions through which those functions and results are to be realized.

Note 1 to entry. The description may consist of one document or a set of linked records, diagrams, and specifications. It may describe how the service establishes the function and result sought during an interaction and how relevant Results and decisions carry forward across interactions.

Note 2 to entry. The description provides a basis for identifying the service and configuration to be examined. Examination establishes how the actual service realizes the applicable requirements under stated conditions.

4. AI Service Quality Object

The AI Service Quality Object Description shall be maintained as an identifiable document or set of linked records. The party responsible for maintaining it shall identify where the records can be retrieved and which version applies to the identified service configuration. In this clause, “the description” refers to the AI Service Quality Object Description defined in 3.5.

NOTE The provider can use existing controlled records to assemble the AI Service Quality Object Description. For example, a service specification can identify the Intended Functions and Intended Results, an architecture record can identify Constituents and Interfaces, and operating procedures can identify who performs and authorizes the Work. The Description can link these records so that the applicable requirements, responsibilities, and configuration can be identified for the service as deployed.

4.1 Intended Functions and Results

4.1.1 The AI Service Quality Object Description shall identify the needs the service is intended to address, its assigned Intended Functions and Intended Results, and the role of AI in performing or supporting those functions. It shall specify how those needs, functions, and results are established, reviewed, and revised across applicable life-cycle stages, including conception, design, deployment, operation, modification, and withdrawal. It shall identify how needs and limitations discovered during operation are considered in subsequent changes to the service.

4.1.2 The description shall identify the functions assigned to the service, as applicable, including interpreting inputs, clarifying needs, generating or transforming content, making recommendations, supporting decisions, carrying out actions, and verifying or following up on results. It shall distinguish the functions assigned to the service from those assigned to recipients, personnel, and connected systems or services.

4.1.3 For each Intended Function, the description shall identify the associated Intended Results, their recipients, and the Interfaces through which those results are delivered or made available. It shall distinguish outputs the service is intended to produce, actions it is intended to complete, and outcomes it is intended to achieve or support.

4.1.4 Where the function to be performed or the result sought for a particular task depends on a recipient's request, the description shall specify how the service establishes that function and result within its assigned role and applicable assistance limits. It shall specify how the recipient can clarify or revise them, what support is necessary for that participation, and who has authority to make or authorize consequential decisions or actions. It shall identify the conditions that must be satisfied before the Work proceeds and what Work may proceed while material uncertainty remains. These provisions also apply where the service establishes the task and performs the Work within a single interaction.

4.1.5 The description shall identify the operating conditions and assistance limits applicable to each assigned function. It shall specify the contributions necessary to perform the function, including, as applicable, recipient input, human review or authorization, and information, action, or verification from a connected system.

4.1.6 Where several requests occur within a service period, the description shall specify how the service distinguishes Work directed toward separate Intended Results from Work that jointly contributes to a shared Intended Result. It shall address later requests that rely on, revise, or supersede earlier Work or Results, whether they occur within one chat, across chats, or through other service Interfaces.

4.1.7 Where Work toward an Intended Result is carried out through one or more interactions, the description shall identify the relevant units of Work, the functions and results assigned to them, and how their Results contribute to the Intended Result of the continuing Work. It shall distinguish the conditions for completion of an individual unit of Work from those for completion of the continuing Work. Where Work extends across interactions or periods of service provision, the description shall specify the intermediate Results, decisions, and unresolved matters that must carry forward, any time limits or permitted interruptions, and the response required when information necessary for resumption is unavailable.

4.1.8 Where achievement of an Intended Result requires review, confirmation, a handoff, or further action after an output is delivered, the description shall identify the remaining Work, the party or system responsible for it, and when that contribution is required.

4.1.9 The description shall link each Intended Function and Intended Result to the relevant Constituents, Work, Interfaces, and external contributions through which it is realized. Where realization depends on a particular model, software, hardware, data, or connected-service configuration, the description shall identify that dependency and the conditions under which it applies.

4.2 Life Cycle, Essential Work, Results, and External Contributions

4.2.1 The description shall identify the life-cycle stages applicable to the AI Service Quality Object and the Essential Work and Results required at each stage to establish, realize, sustain, modify, or end its Intended Functions. It shall identify the Results that must be available before Work proceeds to a subsequent stage.

4.2.2 The description shall distinguish the life cycle of the service from the shorter cycles of Work undertaken for particular requests or continuing tasks. For each type of task, it shall identify, as applicable, the Work required to interpret the request, establish the function and result sought, prepare and deliver an output or action, and carry out any assigned verification or follow-through. These activities may occur within a single interaction or extend across several interactions.

4.2.3 The description shall identify the Essential Work and Results that connect successive stages, tasks, or interactions. Where continuing Work depends on earlier Results, it shall identify which Results must be carried forward, the conditions under which they remain valid, and the Work required when they are unavailable, superseded, or no longer valid.

4.2.4 Where an external party or connected service performs Essential Work or provides a necessary Result, the description shall identify the contribution, the party or service responsible for it, the Interface through which it is provided, and the conditions for accepting and using it. These conditions shall address, as applicable, the required content or action, timing, validity, authorization, and confirmation of delivery or completion. The description shall account for the necessary external Work and Results while identifying the external party or connected service separately in accordance with 4.3.

4.2.5 Where a necessary contribution is delayed, unavailable, incomplete, or invalid, the description shall specify which Work may continue, which Work must be restricted or suspended, and what is required before it resumes. Where Work is transferred to another party or service, it shall distinguish referral or initiation of the transfer from acceptance and completion of the assigned Work.

4.2.6 The description shall identify the Essential Work and Results required when a change to a model, software component, data source, connected tool, or service arrangement can materially affect an Intended Function or Intended Result. It shall specify how the change affects Work already in progress, including which intermediate Results, decisions, or authorizations require re-examination before further reliance.

4.2.7 Where suspension or withdrawal affects unfinished Work or Results on which recipients or connected services continue to rely, the description shall identify the responsibilities for closing or transferring the Work, communicating its status, and making necessary Results available under the applicable conditions.

4.3 Boundaries, Interfaces, and Constituents

4.3.1 The description shall establish the Boundaries of the AI Service Quality Object in relation to its Intended Functions and Intended Results. It shall identify the Constituents and Essential Work and Results within those Boundaries, the external parties and systems on which the service depends, and the Interfaces through which they contribute.

4.3.2 The description shall identify the Constituents that materially affect the realization of each Intended Function or Intended Result. These may include AI models, software that coordinates the Work, data and knowledge sources, records used to carry Work across interactions, connected tools, computing resources, and personnel performing assigned roles. It shall identify the versions, configurations, or selection rules needed to determine which Constituents perform the Work under specified conditions.

4.3.3 Where an external party or connected service performs Essential Work or provides a necessary Result under 4.2.4, the description shall account for that Work or Result within the Quality Object to the extent that it contributes to Function Realization. It shall identify the contributing party or system separately and state the extent to which any of its Constituents are included within the Boundaries. It shall identify where responsibility for providing, accepting, and using the contribution lies.

4.3.4 For each Interface necessary to realize an Intended Function or Intended Result, the description shall identify the parties or Constituents it connects, what passes between them, and the conditions under which the exchange is valid. As applicable, it shall identify the information, instructions, authorizations, actions, status messages, or confirmations exchanged, together with the requirements for timing and the responsibilities on each side.

4.3.5 At Interfaces with recipients, the description shall identify how recipients can supply or correct information, clarify or revise the result sought, and authorize, interrupt, or withdraw an action where applicable. It shall distinguish what the service can observe or verify through the Interface from what depends on a recipient’s statement, an inference, or confirmation from another party or system.

4.3.6 Where the service uses a connected tool or delegates Work to another system, the description shall identify where the action is initiated, where it takes effect, and which party or system confirms its status. It shall distinguish submission, acceptance, execution, and confirmation of the Result wherever those states affect further Work, including recovery after interruption.

4.3.7 Where the selection, substitution, or failure of a model, data source, tool, or connected service changes the capability, responsibility, or assistance limits of the AI service, the description shall identify the configurations in which those changes apply and the Interfaces affected. It shall specify what Work may continue under each configuration and what must be established before normal Work resumes.

4.3.8 The description shall specify how records of actual Work and Results can be linked to the version of the description and the service configuration applicable when the Work was performed. The link shall identify the deployment, server pool, or individual server at the level of detail needed to distinguish differences that materially affect Function Realization. The description shall identify where the records can be retrieved, who is responsible for maintaining them, and the applicable retention periods and the events that start those periods. It shall distinguish the retention of the description and its applicable versions from the retention of conversation, task, and action records.

4.3.9 The AI service provider shall identify in the AI Service Quality Object Description the party accountable for fulfilling each applicable requirement and the party or parties that perform the associated Work. Where different parties are involved, the description shall identify their respective authority, obligations, and Interfaces. The provider shall establish the assignments in accordance with applicable obligations and agreements, confirm acceptance where an assignment depends on another party's agreement, and identify any unresolved gaps. The provider shall maintain these assignments when the service or its external contributions change.

4.3.10 Recurring activities, interactions, and operating conditions may be described by type where the same applicable requirements and responsibilities apply. The Description shall identify variations that materially affect the Work, the Results, or affected persons. Where an applicable clause requires Evidence of a particular occurrence, that Evidence shall be traceable to the applicable Description.

5. Quality Factors and Indicators

The requirements in this clause apply to the AI Service Quality Object identified in accordance with Clause 4. Their applicability shall be determined in relation to the service's assigned Intended Functions and Intended Results and the conditions under which the service undertakes the Work.

5.1 Factor 1 — Object's Functional Requirements

Indicator 1.1 — Completeness of Functional Requirements

5.1.1.1 The Functional Requirements shall be documented in the AI Service Quality Object Description or in identifiable, retrievable records referenced by it. For each assigned Intended Function, they shall state what the AI service must do and what Results it must deliver or support across the applicable life-cycle stages.

5.1.1.2 For each assigned function, the Functional Requirements shall specify the requests or events that initiate the Work, the information and contributions needed to proceed, the Result to be delivered, and the Interface, timing, and conditions for delivery. As applicable, they shall distinguish generating an answer or draft, making a recommendation, preparing a proposed action for approval, and executing or delegating an action. Where the service may move from one of these forms of assistance to another, the requirements shall specify what permits that transition.

5.1.1.3 Where a recipient's request establishes the function or Result sought for a particular task, the Functional Requirements shall specify what the service must establish before undertaking the Work. As applicable, they shall cover the requested deliverable, the item or version to which the request applies, and whether the recipient is asking for a proposal or authorizing execution. They shall specify which Work may proceed while clarification, confirmation, or authorization is pending. An individual recipient request shall not, by itself, constitute a new Functional Requirement of the AI Service.

5.1.1.4 Where an Intended Result depends on Work across multiple interactions, the Functional Requirements shall specify which prior instructions, approved decisions, draft Results, and unresolved matters the service must recover and apply before continuing. They shall specify what the service must do when only part of the earlier context is available or when a saved summary conflicts with a later confirmed instruction. They shall identify which Work may proceed and which Work requires recovery or confirmation of the relevant record.

5.1.1.5 Where the AI service recommends, initiates, or delegates an action whose effects extend beyond the interaction, the Functional Requirements shall specify which party or system authorizes, executes, confirms, and follows up on the action, as applicable. They shall separately identify the required Results for delivery of the recommendation or proposed action, execution of the action, and achievement of the Intended Result.

5.1.1.6 Where the service depends on recipient input, retrieved information, human review, or a connected tool, the Functional Requirements shall specify the contribution needed, when it is needed, and what the service must do if the contribution is delayed, unavailable, incomplete, or invalid. They shall specify which answers, drafts, or other outputs may still be delivered, which actions must be withheld, and the conditions for transferring or resuming the affected Work.

Indicator 1.2 — Clarity and Unambiguity of Functional Intent

5.1.2.1 The documented Functional Requirements shall state each Intended Function and its associated Intended Results in terms that permit examination of whether they are fulfilled. They shall identify the Work assigned to the AI service, the applicable conditions and assistance limits, and the criteria for determining whether the required Result has been delivered or achieved.

5.1.2.2 The Functional Requirements shall distinguish the Intended Functions assigned to the service from the function and result sought for a particular task. Where a recipient's request establishes or clarifies the task, the requirements shall specify how the service distinguishes the recipient's stated purpose, the requested action, and any inference about the result sought.

5.1.2.3 Where different interpretations of a request would materially change the Result, the Work performed, or the consequences of the assistance, the Functional Requirements shall identify what the service must clarify or confirm before proceeding. As applicable, this shall include the referenced item or version, the requested deliverable, and whether the recipient wants a proposal or authorizes execution. Where Work may proceed on a stated assumption, the requirements shall identify the permitted Work and when that assumption must be confirmed before further action.

5.1.2.4 Where the requested action may materially conflict with the recipient's stated purpose, the Functional Requirements shall specify when the service must identify and communicate the conflict, what clarification is needed, and which assistance may proceed while the conflict remains unresolved.

5.1.2.5 Where a later instruction may revise, supersede, or conflict with an earlier instruction or agreed Result, the Functional Requirements shall specify how the service establishes which instruction governs the affected Work and when recipient confirmation is required.

5.1.2.6 The Functional Requirements shall specify what establishes that an answer or draft was delivered, a requested action was completed, or the associated Intended Result was achieved. Where confirmation depends on a recipient, a connected system, or another party, the requirements shall specify what status the service may report while confirmation remains unavailable.

Indicator 1.3 — Explicit Treatment of Functional Boundaries and Interfaces

5.1.3.1 For each Interface necessary to realize an Intended Function or Intended Result, the documented Functional Requirements shall identify the recipient input, retrieved information, review decision, tool action, or other contribution required through that Interface. They shall identify the responsible party or system, when the contribution is needed, and the conditions under which the AI service may use it.

5.1.3.2 Where the AI service receives a request, correction, or revised instruction through a recipient Interface, the Functional Requirements shall specify how it establishes which task, earlier instruction, or draft Result the input applies to. Where applying the input to the wrong item would materially change the Result or action, the service shall resolve the reference before carrying out the affected Work.

5.1.3.3 Where the AI service retrieves information from a connected source to generate an answer, recommendation, or action, the Functional Requirements shall specify which facts require verification, what source and currency are acceptable, and what the service may provide if retrieval fails or sources conflict. Where that uncertainty affects reliance on the Result, the service shall distinguish retrieved or confirmed information from its own inference.

5.1.3.4 Where an AI-generated proposal requires recipient authorization or human review before execution, the Functional Requirements shall specify the action, target, and material terms that must be presented for approval. If the proposed action changes materially after approval, or approval is withdrawn or expires, the service shall withhold execution until applicable approval is established.

5.1.3.5 Where the AI service invokes a connected tool or delegates Work, the Functional Requirements shall specify what must be transmitted and what response establishes acceptance, execution, or confirmation of the Result. If the tool times out after submission and repeating the action could cause harm, the service shall establish the status of the earlier attempt before retrying. Until that status is established, the service shall report it as unresolved.

5.1.3.6 Where a required connected source, tool, or human reviewer becomes unavailable, the Functional Requirements shall specify which outputs or actions the service may still provide and which it shall withhold. Where an alternative is permitted, the requirements shall specify what the service must establish before relying on that alternative for the affected Work.

5.1.3.7 For exchanges on which consequential Work or a reported completion status depends, the Functional Requirements shall identify the information needed to trace the relevant input or approval, the action requested, the response received, and the status reported to the recipient.

Indicator 1.4 — Identification and Documentation of Functional Assumptions and Uncertainties

5.1.4.1 The documented Functional Requirements shall identify the assumptions and material uncertainties on which each assigned function or Intended Result depends. As applicable, they shall address the service's demonstrated capability for the task, the availability and currency of information, the completeness of prior interaction records, and the availability of recipient input, authorization, human review, or connected tools.

5.1.4.2 For each material assumption, the Functional Requirements shall specify the basis for relying on it, the conditions under which it remains valid, and when the service must confirm it. They shall identify which Work may proceed while confirmation is pending and which Work depends on confirmation before it can proceed.

5.1.4.3 Where the AI service uses an inference to fill a gap in recipient instructions, retrieved information, or prior interaction records, the Functional Requirements shall specify when that inference may be used and when it must be communicated to the recipient. Where the inference could materially change a consequential action or its effects, the service shall obtain the confirmation required for that action before proceeding.

5.1.4.4 Where uncertainty affects the reliability or applicability of a generated answer, recommendation, or proposed action, the Functional Requirements shall specify what the service may deliver and what it must withhold. Where qualified assistance is permitted, the service shall identify the unresolved matter and explain how it limits reliance on the Result.

5.1.4.5 Where new information, a change in configuration, or a change in the availability of a necessary contribution undermines an assumption used in earlier Work, the Functional Requirements shall specify when the service must reassess the affected Work. Where the recipient may still rely on an earlier answer, recommendation, or instruction, the service shall identify the affected Result and communicate any correction, qualification, or withdrawal before further consequential Work proceeds.

5.1.4.6 For Work that depends on a material assumption or unresolved uncertainty, the Functional Requirements shall identify the records needed to establish the basis used, any confirmation obtained, and the restriction or qualification applied to the Work or Result.

Indicator 1.5 — Traceability of Functional Requirements Across the Lifecycle

5.1.5.1 Each documented Functional Requirement shall be traceable to the assigned Intended Function or Intended Result that it supports and to the applicable conditions and assistance limits. The records shall identify the version of the requirement applicable to the service configuration under examination.

5.1.5.2 Where a function depends on a model, software component, data source, connected tool, or human contribution, the Functional Requirements shall identify the Work or Result assigned to that Constituent or contributor. Where several Constituents or contributors jointly realize the function, the records shall identify their respective contributions and the Interfaces through which they are coordinated.

5.1.5.3 Where Work continues across interactions, the Functional Requirements shall specify how the service associates the affected Work with the recipient instructions, approved decisions, and versions of earlier Results on which it relies. Where a saved summary is used in place of the underlying exchanges, the requirements shall specify how to establish which decisions and Results the summary represents and whether they still apply.

5.1.5.4 Where the AI service initiates or delegates a consequential action, the Functional Requirements shall specify how to trace the action and its material terms to the applicable authorization, the request sent for execution, the status returned, and the Result reported to the recipient.

5.1.5.5 Where a change to a model, configuration, data source, tool, or service arrangement affects an assigned function, the records shall identify the Functional Requirements affected, the configurations to which the earlier and revised requirements apply, and any Work in progress that must be reassessed before it continues.

Indicator 1.6 — Functional Changes

5.1.6.1 Changes that may affect the Functional Requirements shall be identified, documented, and evaluated before the affected capability is declared available. The evaluation shall address, as applicable, changes to the AI model, instructions that govern its behavior, use of conversation history, information sources, connected tools, permissions, and arrangements for human review.

5.1.6.2 The evaluation shall determine whether each affected Intended Function and Intended Result can still be realized under the specified conditions and assistance limits. It shall identify any required changes to clarification, authorization, execution, verification, or reporting of Results.

5.1.6.3 Where a change may affect continuing Work, the Functional Requirements shall specify which prior instructions, approved decisions, intermediate Results, and authorizations may still be used and which require confirmation or reassessment. The service shall establish the validity of any item requiring reassessment before relying on it in further consequential Work.

5.1.6.4 Where a change reduces an available capability or invalidates a condition on which earlier assistance depended, the Functional Requirements shall specify which Work may continue, which Work must be restricted or suspended, and when the service shall inform affected recipients. Where recipients may still rely on an earlier Result, the requirements shall specify how the service identifies and communicates any necessary correction, qualification, or withdrawal.

5.1.6.5 The records of the change shall identify the Functional Requirements and service configurations affected, the Evidence supporting continued or revised capability, and any restrictions that apply until the affected capability has been verified.

5.2 Factor 2 — Object's Technical Requirements

Indicator 2.1 — Completeness of Technical Specifications

5.2.1.1 The Technical Requirements shall be documented in the AI Service Quality Object Description or in identifiable, retrievable records referenced by it. They shall cover the technical properties and conditions necessary to support the applicable Functional Requirements across the relevant life-cycle stages. The records shall identify the service configurations to which those requirements apply.

5.2.1.2 Where an assigned function depends on an AI model, the Technical Requirements shall specify the model configurations permitted for that function, the supported forms and limits of inputs and outputs, and the required behavior when the service uses a substitute model or reaches a processing limit. Where an output must meet a specified format or satisfy other checkable conditions before further use, the requirements shall specify those conditions and how the service shall handle an output that fails them.

5.2.1.3 Where continuing Work depends on prior interactions or retrieved information, the Technical Requirements shall specify what context the service must be able to obtain and use, and how the applicable versions of instructions, approved decisions, and source records are identified. They shall specify how the service detects when context required for the Work has been omitted, truncated, or replaced by a summary that lacks the required information, and how it limits affected Work until that context is recovered or confirmed.

5.2.1.4 Where an AI-generated output can trigger a consequential action, the Technical Requirements shall specify controls at the point of execution that check whether the action, target, and material terms match the current authorization. They shall specify how execution is withheld when authorization is missing, expired, withdrawn, or inconsistent with the proposed action. Where interruption leaves the status of an earlier attempt unresolved, the requirements shall specify how the service establishes that status or prevents a retry that could duplicate the action.

5.2.1.5 Where the AI service depends on a connected source, tool, or system, the Technical Requirements shall specify the supported exchange formats, the permissions required, and the signals used to distinguish acceptance, execution, failure, and unknown status. They shall specify the technical conditions for detecting a failure, switching to an allowed alternative, and restoring the affected capability.

5.2.1.6 The Technical Requirements shall specify the records needed to examine whether the relevant technical conditions were satisfied during actual Work. As applicable, the records shall allow an examiner to establish which model and service configuration were used, what information or authorization was available, which tool action was requested, what status was returned, and what Result the service delivered or reported.

Indicator 2.2 — Correctness and Unambiguity of Translation from Functional to Technical Requirements

5.2.2.1 The documented Technical Requirements shall identify the Functional Requirements they support and the service configurations to which they apply. Where several models, software components, or connected systems jointly support a Functional Requirement, the records shall identify the technical contribution of each and the conditions required for their combined operation.

5.2.2.2 Where a Functional Requirement calls for the service to use recipient instructions, prior decisions, or retrieved information, the Technical Requirements shall specify how the applicable information is made available to the model or other component performing the Work. They shall address how the service preserves the distinctions required for the function, including which instructions govern the task, which information comes from a source, and which information was generated by the service.

5.2.2.3 Where a Functional Requirement limits the service to proposing an action until authorization is established, the Technical Requirements shall specify how that limit is enforced at the Interface that can execute the action. The control shall apply to each permitted execution path, including a substitute model or connected tool where one is permitted.

5.2.2.4 Where a Functional Requirement calls for the service to report the status of delegated Work, the Technical Requirements shall identify the signals or records that support each status the service may report. A signal confirming submission shall support a report of submission only. A report of execution or achievement of the Intended Result shall require the corresponding Evidence specified for that status.

5.2.2.5 Where model routing, context selection, tool availability, or a change in configuration affects the technical conditions supporting a Functional Requirement, the Technical Requirements shall specify how the service determines whether those conditions remain satisfied. They shall identify the configuration and operating conditions under which the function may continue, proceed with restrictions, or be suspended.

5.2.2.6 The correspondence between Technical and Functional Requirements shall be verifiable from the applicable requirements records, configuration records, and Evidence of service behavior. The records shall enable an examiner to determine whether a specified technical property supports the required function under the conditions in which the service provides it.

Indicator 2.3 — Internal Consistency and Coherence of Technical Specifications

5.2.3.1 The Technical Requirements for AI models, instructions governing model behavior, context selection, connected tools, and execution controls shall be mutually consistent in each permitted service configuration. Where requirements apply at different stages of the same Work, compliance at one stage shall preserve the conditions required at subsequent stages.

5.2.3.2 Where the service combines recipient instructions, prior decisions, retrieved material, or content generated by a model, the Technical Requirements shall establish compatible rules for identifying which material governs the Work and which material may be used as information. Those rules shall apply when context is summarized, shortened, or transferred to another model. Where the status of an instruction or approval cannot be established, the requirements shall specify which Work must be withheld pending resolution.

5.2.3.3 Where a model generates a request for a connected tool, the Technical Requirements for generation, validation, and execution shall use compatible definitions of the action, target, and material parameters. Where an Interface converts values, supplies defaults, or transforms the request, the requirements shall specify how the service verifies that the action to be executed remains within the applicable authorization.

5.2.3.4 The Technical Requirements for timeouts, retries, cancellation, and recovery shall use compatible definitions of action status across the service and its connected systems. Where the service stops waiting for a response but the connected system may still execute the action, the requirements shall preserve that uncertainty in subsequent processing and in the status reported to the recipient.

5.2.3.5 Where the service can switch to a substitute model, source, or tool, the Technical Requirements shall specify the conditions under which the substitute can preserve the applicable context, permissions, validation, and status reporting. Where a required condition cannot be preserved, the requirements shall specify the restriction or suspension that applies to the affected Work.

5.2.3.6 Conflicts among Technical Requirements shall be resolved before the affected configuration is made available for the Work concerned. The resolution shall identify the requirements and configurations affected and be supported by Evidence that the combined service behavior satisfies the applicable Functional Requirements.

Indicator 2.4 — Explicit Treatment of Technical Assumptions and Uncertainties

5.2.4.1 The documented Technical Requirements shall identify the assumptions and material uncertainties affecting the adequacy or applicability of the technical properties specified for each permitted service configuration. As applicable, they shall address variation in model outputs, the completeness of context available to the model, the validity of retrieved information, and the behavior of connected tools and systems.

5.2.4.2 For each material technical assumption, the Technical Requirements shall identify the Evidence that supports it, the configurations and operating conditions to which that Evidence applies, and the conditions that require reassessment. Evidence obtained using one model, instruction set, context limit, or tool version shall be applied to another configuration only where its applicability has been established.

5.2.4.3 Where a technical requirement depends on model outputs meeting specified conditions, the Technical Requirements shall identify which properties can be checked for each output and which can be evaluated only across a defined set of cases or repeated runs. They shall specify the response required when a check fails or when the Evidence is insufficient to establish the required performance. Where an output can initiate a consequential action, the requirements shall specify which properties must be checked before execution.

5.2.4.4 Where the service relies on selected, shortened, or summarized context, the Technical Requirements shall identify the information that must be available to the component performing the Work. They shall specify the checks used to establish its availability and the restrictions that apply when the service cannot establish whether a required instruction, approval, or earlier Result was preserved.

5.2.4.5 Where a technical control depends on the behavior of a connected tool or system, the Technical Requirements shall specify the behavior assumed and how it is verified. As applicable, they shall address whether a timeout can leave an action pending, whether cancellation takes effect before execution, and whether a repeated request can produce a second action. They shall specify the restrictions that apply when the behavior cannot be established.

5.2.4.6 Where a material technical uncertainty remains unresolved, the Technical Requirements shall identify the affected function and configuration, the Work that may continue under specified restrictions, and the Evidence required before those restrictions can be lifted.

Indicator 2.5 — Alignment of Technical Specifications with Lifecycle Conditions

5.2.5.1 For each Interface identified under 4.3 as necessary to realize an Intended Function or Intended Result, the Technical Requirements shall specify the permitted data and operations, the format and meaning of exchanged values, the applicable permissions, and the conditions for accepting or rejecting an exchange. They shall identify which component or connected system performs each required check.

5.2.5.2 Where retrieved material or a tool response is passed to an AI model, the Technical Requirements shall specify how the receiving component identifies its source and permitted use. The service shall prevent instructions embedded in that material from overriding the instructions that govern the Work, granting permissions, or serving as authorization for an action.

5.2.5.3 Where the service transfers context to another model, tool, or external service, the Technical Requirements shall specify which information may cross the Boundary and which must be withheld or transformed before transfer. They shall specify how the receiving component associates the permitted information with the correct recipient, task, and applicable version of the instructions or records.

5.2.5.4 Where the service sends a model-generated action request to a connected tool, the Technical Requirements shall specify the checks that apply to the operation, target, material parameters, and execution permissions at the Interface. Where the request is translated, supplemented with default values, or forwarded to another system, the service shall verify that the action as submitted for execution remains within the applicable authorization.

5.2.5.5 Where the service receives a result or status from a connected system, the Technical Requirements shall specify how it associates the response with the correct action attempt and determines what the response confirms. They shall address late, duplicate, conflicting, and missing responses and specify when the service must withhold dependent Work or report the action status as unresolved.

5.2.5.6 Where an Interface version, permission, or connected system changes, the Technical Requirements shall specify the compatibility checks required before the affected Work proceeds. The records shall allow an examiner to identify the Interface versions used and determine whether a required exchange was accepted, rejected, or left unresolved.

Indicator 2.6 — Traceability of Technical Specifications to Functional Requirements

5.2.6.1 Each documented Technical Requirement shall be traceable to the Functional Requirement it supports, the Constituents and Interfaces responsible for its implementation, and the versions and configurations in which it applies.

5.2.6.2 Where model selection, instructions governing model behavior, context assembly, or tool routing can affect conformity, the Technical Requirements shall specify the records needed to establish which configuration was actually used for the Work examined. The records shall distinguish the configuration approved for use from the configuration that performed the Work, including any substitution or fallback.

5.2.6.3 Where several components jointly implement a technical control, the records shall identify the contribution of each component and the Interface at which the control takes effect. For a model-generated action, the records shall enable an examiner to link the proposed action to the applicable authorization, the request submitted for execution, and the status returned by the connected system.

5.2.6.4 Evidence used to verify a Technical Requirement shall be traceable to the applicable requirement version, the configuration examined, the test or operating conditions, and the observed Result. Where Evidence applies only to specified models, tools, or execution paths, the records shall state that scope.

5.2.6.5 When a Technical Requirement or its implementation changes, the records shall identify the affected Functional Requirements, configurations, and prior verification Evidence. They shall establish which Evidence remains applicable and which technical properties require further verification.

5.2.6.6 Where records needed to establish conformity are missing or cannot be linked to the Work examined, the records of the examination shall identify the link that could not be established and the resulting limit on the Finding.

Indicator 2.7 — Technical Specification Changes

5.2.7.1 Changes to Technical Requirements for the identified AI Service Quality Object shall be identified, evaluated, documented, and authorized before implementation. The evaluation shall address effects on Functional Requirements, related Technical Requirements, assumptions, Boundaries, Interfaces, and applicable life-cycle and operating conditions.

5.2.7.2 Changes to AI models, model versions, instructions governing model behavior, data sources, context assembly, retrieval arrangements, software, tools, tool permissions, external services, security controls, monitoring, and operating configurations shall be evaluated where they may alter the implementation or effectiveness of Technical Requirements.

5.2.7.3 The evaluation shall examine whether the changed configuration preserves required service capability, assistance limits, human authorization and intervention controls, information protection, and protective responses. It shall identify new or altered dependencies, failure modes, and conditions requiring verification or reassessment.

5.2.7.4 Technical Requirements, associated configuration descriptions, Interface specifications, operating restrictions, and Traceability to Functional Requirements shall be updated to reflect the authorized change. The applicable versions and effective conditions shall be identifiable for subsequent Work, Results, and Quality Determinations.

5.2.7.5 Before the changed configuration is relied upon for affected Work, conformity to the applicable Technical Requirements shall be examined to the extent necessary for the change and its consequences. Unverified capability or unresolved conditions shall be identified and subject to appropriate restrictions.

5.2.7.6 Temporary changes, experimental configurations, and rollback arrangements shall have defined authorization, scope, limits, duration or termination conditions, and restoration criteria. Evidence of implementation, examination, and disposition shall be retained.

5.3 Factor 3 — Conformity Assessment

Indicator 3.1 — Defined Basis for Conformity Assessment

5.3.1.1 The responsible party shall document the basis on which conformity of the AI Service Quality Object can be assessed. The basis shall identify the applicable Description, Intended Functions and Intended Results, requirements and acceptance criteria, service configurations and execution paths, operating conditions, and period covered.

5.3.1.2 The responsible party shall identify the Evidence available for the stated scope and disclose material gaps and exclusions. For consequential actions, the Evidence shall distinguish proposal, authorization, submission, execution, and achievement of the Intended Result as applicable.

Indicator 3.2 — Assessment of Conformity to Technical Specifications

5.3.2.1 The responsible party shall retain Evidence of how the implemented service fulfills applicable Technical Requirements under the specified operating conditions. Where material to conformity, the Evidence shall link the applicable model, routing and context, validation controls, tool calls, authorization and execution status to the Work performed.

5.3.2.2 Where the service relies on supplier or component Evidence, the responsible party shall document the configuration and operating conditions to which it applies and its relevance to the integrated service.

Indicator 3.3 — Assessment of Continued Consistency Between Technical and Functional Specifications

5.3.3.1 The responsible party shall maintain traceability between Functional Requirements and the Technical Requirements that realize them. When changes to models, instructions, context, sources, tools, permissions, or human review can affect that relationship, the responsible party shall identify the affected requirements and establish whether the changed service continues to meet them within applicable assistance limits.

5.3.3.2 Where a substituted or unavailable contribution changes what the service can do, the responsible party shall establish the functions that remain supported and apply and communicate the corresponding restrictions.

Indicator 3.4 — Risk-Proportionate Depth of Conformity Assessment

5.3.4.1 The responsible party shall make Evidence available in proportion to the foreseeable consequences of failure. The Evidence shall address material variation across model runs, configurations and execution paths, shared dependencies, and transitions such as fallback, interruption, and handoff. The responsible party shall disclose material limits in coverage.

Indicator 3.5 — Independence and Objectivity of Conformity Assessment

5.3.5.1 For internal and external Conformity Assessment, the responsible party shall make relevant records and configurations accessible under arrangements that protect confidential and personal information and permit independent examination. The responsible party shall disclose material access restrictions, how any supplied cases or records were selected, and any interests or shared dependencies that could affect the independence of the Evidence supplied.

Indicator 3.6 — Treatment of Non-Conformity and Unresolved Conditions

5.3.6.1 For each material nonconformity or unresolved condition, the responsible party shall document the affected Work and Intended Results, the restrictions pending resolution, and the action required to correct the failure or resolve the Evidence gap. Where reliance on an affected Result may continue, the responsible party shall identify and carry out any required correction, qualification, or withdrawal.

5.3.6.2 Where a submitted action may still take effect, the responsible party shall preserve its unresolved status and prevent a repeat attempt that could duplicate it. The responsible party shall retain Evidence supporting resolution before lifting restrictions.

Indicator 3.7 — Traceability of Conformity Statements

5.3.7.1 The responsible party shall ensure that each Conformity Statement it issues or relies on is traceable to the applicable Description, requirements and acceptance criteria, assessed configurations and operating period, and supporting Findings and Evidence. The statement shall identify material exclusions, unresolved conditions, and restrictions on its use.

5.3.7.2 When a change could invalidate the statement, the responsible party shall review its applicability and document whether it was retained, revised, or withdrawn, the basis for that decision, and when the decision took effect.

5.4 Factor 4 — Structural and Functional Integrity

Indicator 4.1 — Definition of Integrity-Critical Functions and Failure Modes

5.4.1.1 The AI Service Quality Object Description shall identify the functions whose loss, degradation, or incorrect execution could materially impair an Intended Result or cause harm. For each such function, it shall identify the conditions that must be maintained for the function to remain available, including any applicable assistance limits.

5.4.1.2 The description shall identify the Constituents, Interfaces, Essential Work, and external contributions on which each integrity-critical function depends. It shall identify shared dependencies whose failure could impair several functions, including a common model, source of governing instructions, context record, authorization control, or connected tool.

5.4.1.3 For each integrity-critical function, the description shall identify foreseeable failure modes and the conditions in which they can arise. As applicable, it shall address loss or misapplication of governing instructions, use of an unverified source or superseded Result, model output that bypasses required validation, execution outside current authorization, duplicate execution, and reporting completion without sufficient confirmation.

5.4.1.4 For each identified failure mode, the requirements shall specify the conditions the service must detect, the latest point at which intervention can still prevent or limit the consequences, and the required response. The response shall identify which Work may continue, which actions must be withheld or suspended, and what must be communicated to affected recipients or connected systems.

5.4.1.5 Where a failure can propagate from generated content to a decision, a tool action, or continuing Work, the requirements shall identify where that propagation must be interrupted. They shall specify the conditions to be verified before the affected function can resume, including the status of any earlier action that may still take effect.

5.4.1.6 The records shall link each integrity-critical function and failure mode to the applicable requirements, configurations, and verification Evidence. They shall identify any failure conditions for which detection, containment, or recovery has yet to be demonstrated and the restrictions that apply.

Indicator 4.2 — Integrity of the Declared Quality Object

5.4.2.1 The AI Service Quality Object shall maintain the structural and functional relationships necessary for its integrity-critical functions under the specified operating conditions. The applicable requirements shall identify the conditions that its Constituents, Interfaces, and external contributions must jointly maintain from initiation of the Work through delivery or confirmation of the Result.

5.4.2.2 Where an integrity-critical function depends on a sequence of AI and non-AI operations, the service shall preserve the association between the task, the governing instructions, the information used, the applicable authorization, and the resulting output or action. Where that association cannot be established, the service shall withhold the affected Work that depends on it.

5.4.2.3 Where generated content can initiate a consequential action, the service shall maintain the required validation and authorization controls through each permitted execution path. Those controls shall remain effective when a request is transformed, forwarded to a connected system, or handled by a substitute model or tool.

5.4.2.4 Where several integrity-critical functions share a model, context record, information source, or execution control, the service shall detect or contain failures of that dependency to the extent specified for the affected functions. A function may continue only where the conditions necessary for its integrity remain satisfied.

5.4.2.5 Where an integrity-critical function is performed jointly with a recipient, human reviewer, or connected system, the service shall establish whether the required contribution was received and remains applicable before relying on it. Where an external action is still pending or its status cannot be established, the service shall preserve that status through subsequent Work and apply the specified restrictions on further action.

5.4.2.6 Evidence of the integrity of the AI Service Quality Object shall address the combined operation of the relevant Constituents, Interfaces, and external contributions under the conditions examined. Evidence from individual components may contribute to the determination where its applicability to the integrated function has been established.

Indicator 4.3 — Consideration of Degradation, Damage, and Aging Mechanisms

5.4.3.1 For each integrity-critical function identified under 5.4.1, the AI Service Quality Object Description shall identify foreseeable mechanisms by which the ability to perform that function can deteriorate over time or through repeated operation. It shall identify the affected Constituents, Interfaces, or external contributions and the conditions under which deterioration could impair an Intended Result.

5.4.3.2 As applicable, the description shall address changes in the types of requests or operating conditions the service encounters, loss of currency or coverage in retrieved information, accumulation of errors in summaries used for continuing Work, and degradation of records needed to recover prior decisions or authorizations. It shall also address loss or corruption of data, configuration records, or other Constituents necessary to maintain the function.

5.4.3.3 The requirements shall specify what observations or checks can reveal deterioration, when they shall be performed, and the criteria for intervention. Where aggregate performance could conceal deterioration in an integrity-critical function, the checks shall address that function under the affected configurations and operating conditions.

5.4.3.4 Where deterioration reaches a specified intervention criterion, the service shall apply the prescribed restriction or suspension to the affected Work. Where prior Results may have been affected, the responsible party shall determine which Results require re-examination and which recipients or connected systems require notice of a correction, qualification, or withdrawal.

5.4.3.5 Before restoring the affected function, the responsible party shall verify that the corrective Work has restored the required performance under the conditions in which deterioration was observed. Where the service depends on reconstructed context or restored records, verification shall establish that the governing instructions, applicable authorizations, and status of pending actions remain available and valid for resumption.

5.4.3.6 The records shall link the detected deterioration to the affected functions, configurations, and operating periods. They shall identify the Evidence supporting the restriction, correction, and restoration decisions and any limitations that remain.

Indicator 4.4 — Integrity Under Foreseeable Abnormal and Transient Conditions

5.4.4.1 For each integrity-critical function, the AI Service Quality Object Description shall identify the foreseeable abnormal or transient conditions that can impair its performance. As applicable, these shall include interrupted responses or tool calls, loss of required context, delayed or conflicting confirmations, temporary loss of authorization services or human review, and transitions between models or tools.

5.4.4.2 For each identified condition, the requirements shall specify what the service must detect, when it must respond, and which integrity conditions must remain satisfied. They shall specify which Work may continue, which Work must be restricted or suspended, and what status shall be communicated to recipients and connected systems.

5.4.4.3 Where an interruption leaves a consequential action unresolved, the service shall preserve the association between the action attempt, its applicable authorization, and any confirmation received. It shall withhold any repeat attempt that could duplicate the action until the earlier attempt is reconciled or duplicate execution is prevented by a verified control.

5.4.4.4 Where the service switches to an alternative model, source, tool, or review arrangement during a disturbance, it shall establish that the alternative preserves the required instructions, context, permissions, validation, and status reporting before relying on it for the affected Work. Where a required condition cannot be maintained, the service shall apply the specified restriction or suspension.

5.4.4.5 Before resuming suspended Work, the service shall verify that the conditions necessary for resumption have been restored. As applicable, it shall reconcile pending actions, confirm that governing instructions and authorizations still apply, and identify any partial or superseded Results that could be mistaken for current or complete Results.

5.4.4.6 Evidence shall demonstrate the required behavior during the disturbance, the transition to restricted operation, and restoration, as applicable. Where the sequence or timing of events can affect the Result, verification shall cover the relevant event orders and timing limits, including confirmations received after a timeout or after recovery begins.

Indicator 4.5 — Preservation of Integrity Through Change and Modification

5.4.5.1 Before a change that could affect an integrity-critical function is introduced, the responsible party shall identify the affected functions and the integrity conditions that must be preserved before, during, and after the change. The evaluation shall cover the affected Constituents, Interfaces, shared dependencies, and external contributions.

5.4.5.2 The evaluation shall address, as applicable, changes to model versions or routing, instructions governing model behavior, context selection or storage, retrieved sources, tool Interfaces, execution controls, and human review arrangements. Where earlier and changed configurations may operate concurrently, it shall identify which combinations can occur and verify that each permitted combination preserves the applicable integrity conditions.

5.4.5.3 Where continuing Work is transferred to a changed configuration, the service shall establish which governing instructions, intermediate Results, authorizations, and pending action records remain valid. It shall preserve the identity and status of actions that may still take effect and withhold any repeat attempt that could duplicate an action until its status is reconciled or a verified control prevents duplicate execution.

5.4.5.4 The requirements shall specify how the change is introduced, the checks required at each stage, and the conditions for proceeding, restricting operation, or suspending the affected function. Where restoration of an earlier configuration is permitted, the service shall verify that the earlier configuration can use the records created during the change and can account for actions initiated under the changed configuration.

5.4.5.5 Before the changed configuration is released for the affected Work, the responsible party shall verify that the integrated service preserves the required relationships among its Constituents, Interfaces, and external contributions. Verification shall cover the transition and the resulting operating conditions, including any restrictions that apply to Work already in progress.

5.4.5.6 The change records shall identify the previous and changed configurations, the integrity conditions examined, the Evidence supporting release or restoration, and any remaining restrictions. The records shall enable an examiner to establish which configuration performed the Work and which conditions applied at the time.

Indicator 4.6 — Integrity Margin and Tolerance to Uncertainty

5.4.6.1 For each integrity-critical function, the requirements shall identify the operating limits within which the service can preserve the conditions necessary for that function. Where the service must intervene before reaching a limit, the requirements shall specify the intervention threshold and justify the margin between that threshold and the limit.

5.4.6.2 Where an integrity margin depends on available context, processing capacity, response time, or the timing of an external contribution, the requirements shall specify how that margin is determined. They shall account for foreseeable variation in model outputs, competing demands on shared resources, and delays in detection, intervention, or confirmation.

5.4.6.3 Where uncertainty affects the integrity of a function, the requirements shall specify how the service determines which Work may continue. If the service cannot establish that the conditions necessary for a consequential action remain satisfied, it shall withhold that action and may provide only the assistance permitted under the remaining conditions.

5.4.6.4 Where the service relies on a model-generated confidence estimate or another indirect indicator to decide whether to continue Work, the requirements shall establish the conditions under which that indicator is valid and the decisions it may support. They shall specify any independent checks required before a consequential action can proceed.

5.4.6.5 Verification shall examine service behavior at and around the specified intervention thresholds and operating limits. Where two or more conditions can deteriorate together, verification shall address combinations that could exhaust the available margin before intervention takes effect.

5.4.6.6 The records shall identify the configurations and conditions for which the integrity margin or tolerance to uncertainty has been demonstrated, the Evidence supporting that determination, and the restrictions that apply beyond the demonstrated conditions.

Indicator 4.7 — Integrity-Relevant Information Retention

5.4.7.1 For each integrity-critical function, the requirements shall identify the information that must be retained to preserve the function during continuing Work, support recovery after interruption or change, and examine a failure. They shall specify when the information must be recorded and how long it must remain available for each of those purposes.

5.4.7.2 The retained information shall enable the service or an examiner, as applicable, to establish which configuration performed the Work, which governing instructions and authorizations applied, which validation checks were performed, and which Results were delivered. Where Work depends on retrieved information or an earlier Result, the records shall identify the source or version used to the extent needed to establish its applicability.

5.4.7.3 For a consequential action initiated through a connected tool, the service shall retain the link between the action attempt, the authorization in force at submission, the request submitted, and any status or confirmation received. Where the outcome remains unresolved, the information needed to reconcile it shall remain available across interruptions and configuration changes until the status is established or responsibility for reconciliation is formally transferred.

5.4.7.4 The requirements shall specify how retained information is protected against loss, unauthorized alteration, and association with the wrong recipient, task, or action. Where records are summarized, migrated, or restored, the service shall verify that the information and links required for the affected function remain retrievable and usable before relying on them.

5.4.7.5 Where information necessary to preserve integrity is missing, damaged, or no longer valid, the service shall apply the specified restrictions to Work that depends on it. Where the information is needed for examination, the Findings shall identify what could not be established and how that limits the determination.

5.4.7.6 The Evidence shall demonstrate that the required information can be retrieved and used within the specified time under the conditions for which it is retained. Where applicable, verification shall cover resumption after interruption, reconciliation of a pending action, and examination of Work performed under a superseded configuration.

5.5 Factor 5 — Harm, Risk, and Safety

Indicator 5.1 — Identification of Hazards, Threats, and Object Vulnerabilities

5.5.1.1 The AI Service Quality Object Description shall identify foreseeable hazards, threats, and vulnerabilities that can arise from the service's outputs, actions, omissions, or interactions. Identification shall cover the relevant life-cycle stages, the conditions of intended use, and reasonably foreseeable misuse.

5.5.1.2 For each identified hazard or threat, the description shall identify who or what could be harmed, the circumstances in which exposure can occur, and the sequence of events through which harm could result. It shall identify the Constituents, Interfaces, or external contributions involved, including those that could allow harm to extend beyond the original interaction.

5.5.1.3 Identification shall consider harm arising both from failure to fulfill an applicable requirement and from foreseeable use of an output or action that fulfills its immediate requirement. As applicable, it shall address plausible but unsupported outputs, misleading statements of certainty or completion, reliance on information that has become invalid, and repeated interactions that can reinforce a harmful course of action.

5.5.1.4 Where the service retrieves information, accepts external content, or invokes connected tools, identification shall address how those Interfaces could be exploited or misused. As applicable, it shall consider unauthorized instructions embedded in retrieved content, disclosure of protected information, changes to an authorized action during transfer, and unintended or duplicate execution.

5.5.1.5 Identification shall address conditions that can increase exposure or reduce the opportunity to detect or prevent harm. As applicable, these shall include the recipient's ability to verify the output, time pressure, absence of effective human support, unavailable external contributions, and shared dependencies whose failure could expose several functions or recipients to the same harm.

5.5.1.6 The responsible party shall document the basis for identification, including the configurations and conditions considered, the Evidence used, and material uncertainties that remain. The record shall be reviewed when changes to the service or its use, incidents, near misses, or recipient reports reveal a new route to harm or materially alter a previously identified one.

Indicator 5.2 — Identification of High-Consequence and Loss-of-Control Scenarios

5.5.2.1 The AI Service Quality Object Description shall identify credible scenarios in which the service could cause or contribute to high-consequence Harm or loss of control. Each scenario shall connect the initiating conditions, relevant vulnerabilities, sequence of events, expected protective responses, and potential consequences for recipients, other affected parties, or connected systems.

5.5.2.2 Where the service provides outputs on which people or systems may rely for consequential decisions, the scenarios shall address how plausible but unsupported content, misleading certainty, or omission of a critical condition could lead to Harm. They shall consider whether repeated interactions, time pressure, or the absence of effective human support could allow reliance to continue before the problem is detected.

5.5.2.3 Where the service can initiate or delegate consequential actions, the scenarios shall address actions that exceed the scope of authorization, affect the wrong target, continue after authorization is withdrawn, or are repeated while an earlier attempt remains unresolved. As applicable, they shall examine whether a stop, cancellation, or restriction reaches each connected system that can still execute the action and takes effect in time to prevent or limit Harm.

5.5.2.4 Scenario identification shall address failures that can propagate across Interfaces or affect several functions or recipients through a shared dependency. As applicable, it shall examine how compromised retrieved content, incorrect context or recipient association, or failure of a shared model, tool, or authorization control could spread the effects of an initial failure.

5.5.2.5 For each scenario, the description shall identify when the service or another responsible party can detect the loss of control, who can intervene, and the latest point at which intervention can still limit the consequences. Where the scenario depends on a human or external system taking protective action, the description shall identify the conditions required for that contribution to be effective.

5.5.2.6 The scenario records shall identify the configurations and conditions examined, the Evidence supporting the scenario, and material uncertainties about its development or consequences. They shall distinguish protective responses assumed in the scenario from responses whose effectiveness has been demonstrated.

Indicator 5.3 — Definition of Risk Control Objectives

5.5.3.1 Risk Control Objectives shall be established for the hazards, threats, and scenarios identified under 5.5.1 and 5.5.2, in accordance with applicable Risk tolerability criteria. Each objective shall specify the protective outcome, who or what is to be protected, the conditions in which protection is required, and the criteria for determining whether the outcome is achieved.

5.5.3.2 Where recipients or connected systems may rely on generated outputs for consequential decisions, the objectives shall specify when a material factual premise must be verified, when unresolved uncertainty must be communicated, and when the service shall qualify or withhold the affected output. Where recipients may continue to rely on an earlier Result that has become invalid, the objectives shall specify the required correction, qualification, or withdrawal and when it must reach those recipients to limit Harm.

5.5.3.3 Where the service initiates or delegates consequential actions, the objectives shall address prevention of execution outside the applicable authorization, execution against the wrong target, and duplicate execution. They shall specify the required outcome when authorization is withdrawn or a stop is requested, including the point by which intervention must take effect and the status the service shall communicate when execution cannot be confirmed.

5.5.3.4 Where protection depends on human intervention or an external contribution, the objectives shall specify what that contribution must accomplish, when it must occur, and what establishes that it has taken effect. They shall specify the protective outcome required while that contribution is delayed, unavailable, or unconfirmed, including which Work may continue under restricted conditions.

5.5.3.5 Where Harm can accumulate across interactions or spread through a shared model, source, tool, or control, the objectives shall specify the extent to which propagation must be prevented or limited. They shall identify the conditions for restricting affected functions and the criteria for restoring them.

5.5.3.6 Where reasonably foreseeable misuse includes deliberately harmful or unauthorized activity, Risk Control Objectives shall specify the assistance the service must withhold or restrict, the conditions under which related assistance may be provided, and the criteria for verifying effective protection. These objectives shall address related requests whose combined Results could enable such activity.

5.5.3.7 The records shall link each Risk Control Objective to the applicable hazard or scenario, service configurations and operating conditions, and acceptance criteria. They shall identify the Evidence needed to determine whether the protective outcome is achieved within the specified time and whether any remaining Risk is within the applicable tolerability criteria.

Informative example. An AI service may support authorized cybersecurity testing while restricting assistance that enables unauthorized intrusion. The assessment considers whether related requests collectively develop an attack plan, even where individual requests appear permissible when examined separately.

Indicator 5.4 — Effectiveness of Risk Controls Across Operating Modes

5.5.4.1 Risk controls shall achieve the protective outcomes established under 5.5.3 in each operating mode and condition for which they are required. Verification shall cover applicable transitions between modes, including interruption, restricted operation, model or tool substitution, and restoration.

5.5.4.2 Detection and protective response shall occur within the time available to prevent or limit the identified Harm. The assessment shall account for delays in detection, communication, decision, execution of the protective action, and any required human intervention. Where protection depends on action by a connected system, the assessment shall establish whether the action takes effect within that time.

5.5.4.3 Risk controls required during a disturbance shall remain effective under the conditions they are intended to control. Where a control depends on the same model, context record, tool, communication path, or human review arrangement that may be impaired, the service shall demonstrate that the control remains effective or provide an alternative that achieves the required protective outcome.

5.5.4.4 Protective actions shall be coordinated across the service and affected connected systems. Where the service restricts an execution path, withdraws an authorization, or requests cancellation, it shall establish which pending actions can still take effect and maintain the required restrictions until their status is resolved. The assessment shall examine whether the protective action introduces or transfers an unaddressed hazard to another function, recipient, or connected system.

5.5.4.5 Where a required control is unavailable, bypassed, or impaired, affected Work may continue only within specified restrictions and where any required compensating control has been verified under the prevailing conditions. Before normal operation resumes, the service shall verify that the required controls are effective and address any pending actions or prior Results that remain relevant to protection.

5.5.4.6 Evidence shall show whether the required protective outcome was achieved under the modes, disturbances, and transitions examined. Where timing or event order can affect protection, verification shall include relevant sequences such as authorization being withdrawn after submission, a confirmation arriving after a timeout, or a substitute model or tool taking over while a restriction is active.

Indicator 5.5 — Systemic and Cumulative Risk Consideration

5.5.5.1 Risk arising from interactions among Constituents, shared dependencies, and cumulative effects shall be identified and evaluated for the integrated AI Service Quality Object. The evaluation shall cover effects that can emerge across successive interactions, delegated actions, or uses of a common Result.

5.5.5.2 The evaluation shall identify conditions that could impair several functions or protective measures at once. As applicable, it shall examine shared models, sources of governing instructions, retrieved data, context records, tool services, authorization controls, and human review arrangements.

5.5.5.3 Where protection depends on separate or duplicate models, tools, checks, or reviewers, the evaluation shall establish whether a common dependency or exposure could defeat them together. As applicable, it shall examine whether a substitute model or an AI-based reviewer uses the same compromised source, incomplete context, or decision criteria as the component it replaces or checks.

5.5.5.4 The evaluation shall trace how an error, unauthorized instruction, or incomplete Result can propagate through context summaries, shared records, delegated Work, or connected tools. It shall identify where the service must interrupt propagation, which downstream Results or actions require re-examination, and which affected recipients or systems require notice.

5.5.5.5 Where repeated outputs or actions can accumulate Harm, the requirements shall specify which effects must be considered together, over what period or sequence of events, and under what conditions intervention is required. As applicable, they shall address repeated reliance on an unsupported premise, successive actions that individually fall within authorization but collectively exceed an applicable limit, and reuse of an erroneous Result across several tasks.

5.5.5.6 The assessment shall examine combined failures and sequences of Work where they can materially affect a Risk Control Objective. The Findings shall identify the configurations, shared dependencies, and cumulative effects examined, the Evidence supporting the determination, and any remaining uncertainty about the extent of Harm or the effectiveness of the controls.

Indicator 5.6 — Preservation of Risk Mitigation Alignment Under Change

5.5.6.1 Before implementing a planned change, the responsible party shall evaluate its effects on the hazards and scenarios identified under 5.5.1 and 5.5.2, the Risk Control Objectives established under 5.5.3, and the effectiveness of the controls and residual Risk. When an unplanned change is detected, the responsible party shall promptly evaluate its effects and restrict affected Work where continued protection cannot be established.

5.5.6.2 The evaluation shall address changes to model versions or routing, instructions governing model behavior, context selection, information sources, tool capabilities, permissions, and arrangements for human intervention, as applicable. Where a change introduces a new route to Harm, exposes additional recipients or systems, or reduces the time available to intervene, the responsible party shall update the affected scenarios, objectives, and controls before permitting the affected Work.

5.5.6.3 Where earlier and changed configurations operate concurrently during testing, staged release, or restoration, the evaluation shall examine protection in each permitted combination. As applicable, it shall verify that restrictions remain effective when Work moves between configurations, that authorization still applies to the action as executed, and that a stop or cancellation can reach any connected system that may still act.

5.5.6.4 Before releasing the changed configuration for the affected Work, the responsible party shall verify that the required controls achieve their protective outcomes within the specified time. Verification shall cover affected execution paths, shared dependencies, and any compensating controls used during the change. Where the required outcome cannot be demonstrated, the affected Work shall remain restricted or suspended.

5.5.6.5 Where the change affects Work already in progress or Results on which recipients or connected systems may still rely, the responsible party shall determine which pending actions must be reconciled and which Results require re-examination. The requirements shall specify when to communicate a correction, qualification, or withdrawal and what protection shall remain in effect pending resolution.

5.5.6.6 The change records shall identify the configurations and operating conditions affected, the hazards and Risk Control Objectives reviewed, the Evidence supporting control effectiveness, and the basis for determining whether residual Risk meets the applicable tolerability criteria. They shall identify any remaining restrictions and when the release, suspension, or restoration decision took effect.

Indicator 5.7 — Risk Information Retention and Transparency

5.5.7.1 The responsible party shall retain the information needed to establish and examine identified hazards, credible scenarios, Risk Control Objectives, control effectiveness, and residual Risk. The records shall link that information to the applicable AI Service Quality Object Description, the service configuration and execution paths concerned, and the conditions and period to which it applies.

5.5.7.2 The records shall include, as applicable, risk assessments, acceptance criteria, control settings, verification results, incidents, near misses, overrides, corrective actions, and changes affecting protection. Where an output or delegated action contributed to an incident or near miss, the records shall enable an examiner to establish which configuration performed the Work, which controls applied, what was communicated to affected parties, and what outcome could be confirmed.

5.5.7.3 The service shall make current restrictions, impaired or bypassed controls, temporary protective measures, and unresolved action status available to those responsible for operating the service, authorizing Work, or intervening. The information shall identify the affected functions and the conditions under which those functions may continue or resume.

5.5.7.4 Where recipients or connected parties must make a decision or take protective action, the service shall communicate the material limitations, uncertainties, and restrictions that affect that decision or action. The information shall distinguish confirmed conditions from unresolved matters and shall be available in time for the intended party to act.

5.5.7.5 The requirements shall specify the retention periods, access conditions, and protections for risk records. They shall define what information is necessary to reconstruct a material event or verify a protective response, including how to preserve the relevant links when recipient content is redacted, summarized, or deleted.

5.5.7.6 The Evidence shall demonstrate that the relevant information can be retrieved and understood by the intended recipient, operator, or examiner within the time required for its use. Where a record is incomplete or unavailable, the Findings shall identify which protective outcome or conclusion about residual Risk could not be established.

5.6 Factor 6 — Object and Environment

Indicator 6.1 — Identification of Bidirectional Environmental Interactions and Exposure Pathways

5.6.1.1 The AI Service Quality Object Description shall identify the conditions outside the Boundary of the AI service that can materially affect realization of its Intended Functions or Intended Results. It shall also identify the material effects that the service’s operation, outputs, or actions can have outside that Boundary. Using the Boundary established under 4.3, the Description shall identify the external parties, systems, and resources that can materially affect an Intended Function or Intended Result, or that can be materially affected by the service’s operation, outputs, or actions. Depending on the service, these may include external organizations, information sources, receiving systems, human support arrangements, and shared computing or physical infrastructure.

5.6.1.2 For each material interaction or recurring interaction type described under 4.3.10, the description shall identify the external condition or affected party, system, or resource, the direction and pathway of influence, and the operating conditions under which the interaction occurs. Where the interaction occurs through an Interface, the description shall identify that Interface and the Constituents or external contributions involved.

5.6.1.3 The description shall identify the external conditions on which the service relies. As applicable, these shall include the languages and input formats encountered in use, the availability and currency of external information, network and computing availability, and the availability of human or institutional support. It shall identify where variation in those conditions can change the service's capability, assistance limits, or the validity of a Result. Where conditions in the recipient's environment can materially affect the suitability, timing, or safety of assistance, the description shall identify the relevant conditions, how they can become known during use, and how the service shall adapt when they change or remain unknown. This includes, where material, whether timely human support is available and whether the recipient can use the proposed next step under the circumstances. The service shall seek clarification when the answer could materially change the assistance and the delay is acceptable. Where clarification is unavailable or delay would be consequential, the service shall apply the assistance limits appropriate to the unresolved condition.

5.6.1.4 The description shall identify how the service's outputs and actions can change the environment in which later Work is performed. As applicable, it shall address generated content entering shared records or information sources, delegated actions changing the state of connected systems, and decisions based on the service's outputs changing the information or conditions encountered in subsequent interactions.

5.6.1.5 Where operation of the service places material demands on external computing or physical resources, the description shall identify those demands and the basis for attributing them to the identified service. As applicable, it shall address computing capacity, electricity, cooling, and water use, including where the service depends on shared or externally operated infrastructure. Where direct measurement is unavailable, the description shall identify the allocation or estimation method and its limitations.

5.6.1.6 The records shall link each material interaction to the affected functions, Results, configurations, and operating conditions. They shall distinguish observed conditions and effects from those established by estimation or assumption and identify any material uncertainties that remain.

Indicator 6.2 — Prevention and Limitation of Unintended Releases and Environmental Harm

5.6.2.1 For material outward interactions identified under 5.6.1, the responsible party shall specify the conditions for authorized release of generated content, recipient information, and requests for external actions. The conditions shall identify the permitted destination, audience, purpose, and extent of release, and any approval required before release.

5.6.2.2 The AI service shall apply those conditions to the destination, audience, and content as they will be presented to the receiving system. Where a tool or connected service expands the audience, adds data from context, or changes the destination, the AI service shall withhold release unless the resulting transmission meets the applicable conditions.

5.6.2.3 Where generated outputs enter shared records or downstream automation, the responsible party shall establish limits on propagation and specify the conditions that require containment. Where an unintended release is detected, the service shall restrict further propagation within its control and communicate the affected destination, known effects, unresolved status, and available corrective actions to the parties responsible for containment.

5.6.2.4 Where AI Work can cause material harm through abnormal demand on shared infrastructure or physical resources, the responsible party shall specify the applicable operating limits and the response when those limits are reached or cannot be confirmed. The response shall address AI-specific sources of excess demand where applicable, including repeated tool calls, retries, and fallback to more resource-intensive execution.

5.6.2.5 The responsible party shall retain Evidence linking a material unintended release or exceedance to the Work, configuration, affected external systems or resources, and containment actions. The records shall distinguish confirmed containment from effects that remain outside the service's control or unresolved.

Indicator 6.3 — Environmental Interactions Under Abnormal and Degraded Conditions

5.6.3.1 The responsible party shall identify abnormal and degraded external conditions that can materially change the interactions described under 5.6.1. These shall include, where applicable, stale or conflicting external data, loss of human review, partial tool outages, delayed acknowledgments, and constraints on shared infrastructure.

5.6.3.2 For each material condition, the responsible party shall specify how the service recognizes the condition or determines that a required external condition cannot be confirmed. The specification shall identify the affected functions, applicable assistance limits, and the criteria for continuing, restricting, suspending, or handing off the Work.

5.6.3.3 The AI service shall apply the specified response within the time required to limit adverse effects. Where a connected system may already have acted, the service shall preserve the unresolved status and withhold any repeat action that could duplicate or conflict with the first action until its status is reconciled.

5.6.3.4 Before restoring affected functions, the responsible party shall establish that the required external conditions have been restored or that an authorized substitute supports the Intended Functions within the applicable limits. The service shall reconcile pending external actions and reassess material information, approvals, or assumptions that may have changed during the disruption.

5.6.3.5 The responsible party shall retain Evidence identifying the observed condition or uncertainty, the affected Work and external Interfaces, the timing and effects of restrictions or handoffs, and the basis for restoration. The records shall distinguish confirmed external effects from effects that remain unresolved.

Indicator 6.4 — Consideration of Cumulative and Long-Term Environmental Effects

5.6.4.1 The responsible party shall identify material effects outside the AI Service Quality Object that can accumulate across interactions, deployments, or recipients, or persist after the Work ends. The analysis shall address, where applicable, repeated AI Work on shared computing resources, persistent changes to external records, and generated material that re-enters external information sources used by the service or by other systems.

5.6.4.2 For each material effect, the responsible party shall specify the relevant time horizon, affected population, system, or resource, and the basis for attributing the effect to the service. The analysis shall address expected and reasonably foreseeable levels of use, including changes in model routing, tool use, or retention that can materially alter the effect.

5.6.4.3 The responsible party shall specify measurable criteria for determining when accumulated or persistent effects require action, how those effects shall be tracked, and what response shall be taken. Where direct measurement is unavailable, the responsible party shall document the estimation method, its material assumptions and uncertainties, and how those uncertainties affect the response.

5.6.4.4 The responsible party shall reassess the effects when use, configuration, or external conditions change materially or when Evidence indicates that prior estimates no longer apply. Where material effects can persist after the service is changed or withdrawn, the responsible party shall specify who is responsible for any continuing monitoring or corrective action, and when that responsibility ends.

Indicator 6.5 — Environmental Protection Alignment Under Change

5.6.5.1 Before implementing a change that can materially alter an interaction under 5.6.1 or an effect under 5.6.2 through 5.6.4, the responsible party shall determine whether the environmental protection requirements and controls remain adequate for the changed service. The determination shall address, as applicable, changes in model routing, tool permissions, downstream use, retention, workload, and hosting arrangements.

5.6.5.2 Where the change alters an affected party, external system, resource, or pathway of influence, the responsible party shall update the Description and the applicable limits, response criteria, and assignments of responsibility before the changed function is used. Where old and new configurations operate concurrently, the responsible party shall account for their combined effects and any differences in the controls applicable to each.

5.6.5.3 The responsible party shall establish the conditions for releasing, restricting, or reversing a change and the criteria for determining whether protective controls remain effective during and after the transition. Where reversal cannot undo an external effect, the responsible party shall specify the containment or corrective action available and who can perform it.

5.6.5.4 The responsible party shall retain Evidence linking the change and affected configurations to the environmental effects considered, revised controls and assignments, release decision, and observations following implementation. Where observations exceed specified limits or invalidate a material assumption, the responsible party shall apply the specified response and reassess continued operation.

Indicator 6.6 — Understanding and Limitation of Emissions and Discharges

5.6.6.1 The responsible party shall identify material physical emissions and discharges attributable to the provision of the AI service, including those arising through external contributions. The Description shall distinguish releases at facilities used to provide the service from emissions associated with purchased energy and shall state the activities, facilities, and periods included in each determination.

5.6.6.2 For each material emission or discharge, the responsible party shall specify the quantity or rate to be determined, the applicable limits or control criteria, and the conditions under which they apply. Where infrastructure is shared, the responsible party shall document the basis for allocating a share to the AI service, the sources and dates of the underlying data, and the uncertainty of the allocation. Estimated shares shall be distinguishable from measured facility totals.

5.6.6.3 The responsible party shall specify and implement measures within its authority to limit attributable emissions and discharges. Where controls depend on an external operator, the Description shall identify who operates the controls and what information the responsible party needs to determine whether they are effective. The responsible party shall specify the response when a limit is exceeded or when missing or uncertain information prevents determining whether a limit is met.

5.6.6.4 The responsible party shall retain Evidence linking the determination and any corrective action to the relevant AI Work, infrastructure, operating conditions, and period. Where workload routing, energy supply, or cooling arrangements change materially, the responsible party shall review whether the prior determination remains applicable in accordance with 5.6.5.

Indicator 6.7 — Environmental Information Retention and Transparency

5.6.7.1 The responsible party shall retain the environmental information and supporting Evidence needed to reconstruct determinations under 5.6.1 through 5.6.6. The records shall link material effects, incidents, and corrective actions to the applicable Description, AI Work, deployment or external contribution, operating conditions, and period.

5.6.7.2 Where a determination relies on shared infrastructure or supplier information, the records shall identify the source, date, scope, and version of the information, together with the allocation method, assumptions, and material uncertainties. The records shall distinguish measurements, estimates, and effects that could not be determined. Changes or corrections shall be traceable to the original determination and to affected Quality Claims.

5.6.7.3 The responsible party shall specify the retention period, the event that starts it, and the arrangements for authorized retrieval. Retention shall account for applicable obligations, the period over which material effects may persist, and the Evidence needed for follow-up and assessment. Where essential records are held externally, the responsible party shall identify who holds them and how authorized access will be maintained after changes in deployment or supplier.

5.6.7.4 The responsible party shall make environmental information available to intended recipients and authorized assessors at a level of detail appropriate to their purpose and access rights. Reported results shall state their scope, period, and material exclusions and uncertainties. Where access restrictions or unavailable information limit a determination, the responsible party shall disclose the resulting limitation on any associated Quality Claim.

5.7 Factor 7 — Functional Reliability

Indicator 7.1 — Identification of Reliability-Critical Functions

5.7.1.1 The responsible party shall identify the Intended Functions for which interruption, delay, intermittent failure, or inconsistent performance can materially prevent realization of an Intended Result. For each function, the Description shall identify the affected Results and recipients, the operating conditions, and the time interval over which dependable performance is required.

5.7.1.2 The identification shall account for Essential Work across interactions, sessions, and handoffs where continuity is necessary. It shall identify dependencies whose loss or intermittent operation can prevent completion, including model routing, retained context, authorization, connected tools, and human review where applicable.

5.7.1.3 The responsible party shall identify how a failure of each reliability-critical function would become observable at the service boundary, including cases in which the service remains available but delivers a late, incomplete, inconsistent, or duplicate Result. The Description shall record the basis for treating each function as reliability-critical.

5.7.1.4 The responsible party shall review the identification when an Intended Function, reliance by recipients, or a material dependency changes, and when operational Evidence reveals a previously unrecognized failure that can materially affect an Intended Result.

Indicator 7.2 — Consistency of Functional Performance Under Expected Operating Conditions

5.7.2.1 For each reliability-critical function identified under 5.7.1, the responsible party shall specify the expected operating conditions and measurable criteria for consistent realization of the Intended Result. The criteria shall define what constitutes an acceptable Result, the time allowed for its delivery where material, and the permitted frequency of failure over a specified observation period or number of opportunities.

5.7.2.2 The criteria shall distinguish permissible variation in generated outputs from variation that changes a required fact, decision, action, or other property material to the Intended Result. Where the service acts through a connected tool, the criteria shall account for the Result confirmed at the receiving Interface and for late, incomplete, or duplicate actions.

5.7.2.3 The responsible party shall demonstrate that the criteria are met across the expected operating conditions. The demonstration shall account for relevant variation in inputs, workload, interaction length, model routing, and connected dependencies. Where overall results can conceal materially poorer performance in a configuration or operating condition, the responsible party shall determine performance for that configuration or condition separately.

5.7.2.4 The responsible party shall retain Evidence of observed performance, including the applicable configuration, operating conditions, observation period, and number of eligible opportunities. The Evidence shall distinguish successful, failed, and unresolved outcomes and identify material limits on the determination.

Indicator 7.3 — Tolerance to Operational Variability and Disturbance

5.7.3.1 For each reliability-critical function, the responsible party shall specify the disturbances the service is required to tolerate, their applicable magnitude or duration, and the required performance during and after each disturbance. Where foreseeable disturbances can occur together and materially affect performance, the required tolerance shall address their combined effects.

5.7.3.2 The service shall maintain the specified performance within those limits. Where a disturbance exceeds the limits or the service cannot establish that required conditions remain satisfied, the service shall apply the specified restrictions and communicate the affected capabilities and status of pending Work to the relevant recipients or connected systems.

5.7.3.3 Where the service uses retries, fallback, rerouting, or handoff to tolerate a disturbance, it shall preserve the context and authorization necessary for the affected Work and reconcile the status of any external action before retrying where duplication would be material.

5.7.3.4 The responsible party shall retain Evidence of performance under the specified disturbances, including their magnitude and duration, the configuration examined, the Results achieved, and any limits on the determination.

Indicator 7.4 — Recovery Capability Following Disruption

5.7.4.1 For each reliability-critical function, the responsible party shall specify recovery requirements for disruptions that prevent the required performance. The requirements shall identify the event from which recovery time is measured, the maximum allowed time where material, the acceptable loss or staleness of Work state, and the conditions that demonstrate restoration of the function.

5.7.4.2 Before resuming interrupted Work, the service shall establish the status of completed, pending, and unresolved steps, and determine whether restored context, instructions, authorizations, and external state remain valid. Where the service cannot establish a prerequisite for resumption, it shall restrict the affected Work and obtain confirmation or use the specified handoff or restart procedure.

5.7.4.3 The service shall reconcile material external actions before resuming or repeating Work where unresolved effects could cause duplication or conflict. It shall communicate to affected recipients or connected systems what has been restored, what remains unresolved, and any action required of them.

5.7.4.4 The responsible party shall retain Evidence that the specified recovery capability is demonstrated for the applicable disruptions and configurations. The Evidence shall identify the disruption and recovery events, the elapsed time, the state recovered or lost, and the Results achieved after resumption.

Indicator 7.5 — Reliability Implications of Degradation and Aging

5.7.5.1 For each reliability-critical function, the responsible party shall identify foreseeable deterioration over time or with use that can materially reduce the service's ability to meet the criteria in 5.7.2. The identification shall address, where applicable, stale source information, loss of compatibility with external dependencies, accumulated Work state, and declining performance as operating conditions diverge from those previously examined.

5.7.5.2 The responsible party shall specify the conditions or trends that trigger review or intervention, the maximum interval between checks where material, and the action required when the criteria for consistent performance are threatened or no longer met.

5.7.5.3 Where deterioration can materially affect an Intended Result, the service shall apply the specified restrictions or corrective action. Before removing restrictions, the responsible party shall establish that the affected function meets the applicable performance criteria under the conditions in which it will resume.

5.7.5.4 The responsible party shall retain Evidence linking observed deterioration and interventions to the affected function, configurations, observation periods, and Results. Where configurations or evaluation methods have changed, the Evidence shall identify material limits on comparisons over time.

Indicator 7.6 — Reliability Under Change and Modification

5.7.6.1 Before introducing a change that can materially affect a reliability-critical function, the responsible party shall identify the affected Intended Results, operating conditions, dependencies, and applicable performance criteria. The review shall cover changes to models, routing, instructions, tools, and external contributions where applicable.

5.7.6.2 The responsible party shall establish that the changed configuration meets the applicable criteria under 5.7.2 and, where affected, the tolerance and recovery requirements under 5.7.3 and 5.7.4. Where prior Evidence is used, the responsible party shall document why it remains applicable and address any conditions or interactions it does not cover.

5.7.6.3 Where Work can span the change or be performed by overlapping configurations, the service shall preserve the context, authorizations, and action status needed for the Intended Result. The responsible party shall specify the conditions for continuing, restricting, or reversing the rollout and how unresolved Work and external actions will be handled.

5.7.6.4 The responsible party shall retain Evidence linking the change, applicable criteria, release decision, and post-change performance. Where observations show that the criteria are no longer met, the responsible party shall apply the specified response and verify restoration before lifting restrictions.

Indicator 7.7 — Reliability-Relevant Information Retention

5.7.7.1 The responsible party shall retain the specifications and Evidence needed to reconstruct determinations under 5.7.1 through 5.7.6. The retained information shall link the applicable performance criteria and their versions to the function, service configuration, operating conditions, observation period, and Results examined.

5.7.7.2 The retained Evidence shall allow authorized parties to reconstruct how eligible opportunities and successful, failed, and unresolved outcomes were counted. Where relevant, it shall link disturbances, recovery events, interventions, and changes to affected Work and Results. Where Evidence is aggregated or redacted, the responsible party shall preserve the information needed to substantiate the determination and document any resulting limits.

5.7.7.3 The responsible party shall specify retention periods, the events that start them, and arrangements for authorized retrieval, including for Evidence held by external contributors. The arrangements shall preserve traceability across material changes to the service, its dependencies, or the record systems, and protect confidential and personal information.

5.7.7.4 The responsible party shall make sufficient retained Evidence available to authorized assessors under appropriate access controls. Where Evidence is unavailable or insufficient to substantiate a reliability determination, the responsible party shall document the gap and its effect on any associated Quality Claim.

5.8 Factor 8 — Human and Object Interaction

Indicator 8.1 — Identification of Human–Object Interactions

5.8.1.1 The responsible party shall identify human interactions with the AI service that materially affect realization of Intended Functions and Intended Results or the persons affected by the Work. The Description shall identify the relevant human roles, including recipients, operators, reviewers, and external parties where applicable.

5.8.1.2 For each material interaction or recurring interaction type described under 4.3.10, the Description shall identify the initiating event, the information or action expected from each party, the applicable Interface, and how the interaction affects the Work or Result. It shall distinguish actions initiated by a person from those initiated by the service and identify where acknowledgment, authorization, correction, intervention, or handoff is required.

5.8.1.3 The responsible party shall identify foreseeable conditions in which an interaction may be absent, delayed, ambiguous, or misunderstood, and specify how the service shall respond where those conditions materially affect the Work or Result. Where a human decision is required, the service shall represent its status accurately and apply the applicable assistance limits while the decision remains unresolved.

5.8.1.4 The responsible party shall review the identification when human roles, service autonomy, or Interfaces change materially, or when operational Evidence reveals a previously unidentified interaction.

Indicator 8.2 — Suitability of Object for Human Interaction

5.8.2.1 For the material interactions identified under 5.8.1, the responsible party shall specify the characteristics of the Interface needed for each human role to perform its expected actions under the intended conditions of use. The requirements shall address, where material, language, accessibility needs, available time, and the consequences of misunderstanding or delay. Recurring interactions may be described by type under 4.3.10.

5.8.2.2 The service shall present the information needed for a person to act in a form that the intended recipient can perceive and understand. Where material to the decision, it shall distinguish a proposed action from an executed action, communicate uncertainty and applicable assistance limits, and identify what decision or action is required from the person.

5.8.2.3 Where a person is expected to review, correct, authorize, stop, or take over Work, the service shall provide an Interface through which that person can perform the action in time to affect the Result. Where the required interaction cannot be completed through the available Interface, the service shall apply the specified alternative, handoff, or restriction.

5.8.2.4 The responsible party shall establish criteria for determining whether intended persons can complete the required interactions and shall retain Evidence of performance under representative conditions of use. The Evidence shall identify the Interface version and material limits on the conditions and persons represented.

Indicator 8.3 — Prevention of Unsafe Human–Object Interaction

5.8.3.1 For material human interactions identified under 5.8.1, the responsible party shall assess reasonably foreseeable pathways to Harm under intended conditions of use and reasonably foreseeable misuse. As applicable, the assessment shall address reliance on generated recommendations, transition from advice to execution, correction of information after Work begins, and human intervention before a consequential action. The responsible party shall involve relevant technical, domain, and human factors expertise and draw on Evidence from intended users or their representatives where needed. Interaction types may be described under 4.3.10.

5.8.3.2 Where acting on an unresolved assumption in a generated recommendation could lead to Harm, the service shall identify the assumption to the recipient before presenting an actionable recommendation that depends on it. If the assumption remains unresolved, the service shall restrict its assistance to the specified limits.

5.8.3.3 Where human authorization is required, the service shall present the proposed action, target, scope, and material consequences before requesting authorization. It shall bind authorization to the action presented and obtain renewed authorization before executing a materially changed action. Silence, timeout, or ambiguous input shall not satisfy the authorization requirement.

5.8.3.4 When a person corrects information on which pending Work depends, the service shall identify the affected steps and invalidate or revise them before continuing. It shall report any affected action already initiated that may still take effect.

5.8.3.5 Where prevention depends on human intervention, the service shall provide the necessary information and controls before the last effective point. If the intervention cannot be obtained in time, the service shall apply the specified restriction, suspension, or handoff. Upon an authorized stop request, it shall prevent further affected Work within its control and communicate the status of actions already initiated and any unresolved external effects.

5.8.3.6 For the identified AI service, the responsible party shall specify in accordance with Section 6 the triggers, assistance limits, latest effective intervention points, and measurable acceptance criteria applicable to the interaction types assessed. It shall retain Evidence of preventive performance under representative conditions, including missing, ambiguous, or withdrawn human decisions where applicable.

Indicator 8.4 — Safety of Human Interaction Under Abnormal Conditions

5.8.4.1 For the material human interaction types identified under 5.8.1, the responsible party shall document in the AI Service Quality Object Description, or in identifiable, retrievable records referenced by it, the abnormal conditions that can impair a person's ability to understand, authorize, correct, stop, or take over Work. As applicable, these shall include interrupted output, missing or stale conversation context, failure to deliver an intervention request or obtain the required human response, conflicting reports about an external action, and fallback configurations that lack required controls.

For each applicable condition, the Description shall identify

a) the affected human interaction, Work, and Intended Result

b) the observable signals or records, their sources, and the criteria the service uses to determine that the abnormal condition exists

c) the evidence required to establish that a prerequisite for continuing Work is satisfied, and when missing, late, stale, or conflicting evidence requires the service to treat that prerequisite as unconfirmed

d) the protective response for a confirmed abnormal condition or an unconfirmed prerequisite, including what the service communicates to the affected person

e) the latest effective point by which the response must take effect.

Recurring conditions may be described by type in accordance with 4.3.10. The Description shall identify variations that materially change detection, the required response, or the time available to act.

5.8.4.2 When interrupted or degraded delivery can leave a person with actionable but incomplete guidance, the service shall identify the affected guidance as incomplete or superseded at the next available contact. Where a qualification, prerequisite, or warning is necessary to prevent Harm, the service shall present it before or with the actionable content. Before resuming consequential Work after context loss, the service shall verify the material instructions, corrections, and authorization on which that Work depends or obtain them again.

5.8.4.3 Where a required human decision or intervention cannot be obtained through the primary Interface, the service shall use an authorized alternative only if it can preserve the required authority, information, and time to act. Otherwise, the service shall restrict or suspend affected Work before the last effective point. A sent notification or a delivery receipt shall not satisfy a requirement for a human decision or intervention.

5.8.4.4 Where a timeout, interruption, or contradictory tool response leaves an external action unresolved, the service shall distinguish the last confirmed state from the unconfirmed outcome and communicate what the person can do next. Where a repeated action could duplicate or conflict with the original, the service shall prevent re-execution within its control until reconciliation and warn against repetition through another channel while the outcome remains unresolved.

5.8.4.5 Before resuming affected Work, the service shall confirm that the required human controls are available and effective, reconcile pending actions, and obtain renewed authorization where the approved action or its material conditions have changed. The responsible party shall specify acceptance criteria and retain Evidence of the trigger, response timing, information presented to the person, action status, and basis for resumption.

Indicator 8.5 — Clarity of Object for Its Users

5.8.5.1 The responsible party shall specify in the Description, or in retrievable records referenced by it, what users need to understand about the service to request Work, interpret Results, and decide whether to rely on them. The specification shall identify what must be communicated, to whom, at which Interface and point in the interaction, and which changes require renewed communication. Recurring interaction types may be described under 4.3.10.

5.8.5.2 At the point where the distinction affects a user's decision, the service shall make clear whether it can provide information, prepare a draft, request a human decision, or execute an action. Where an action is involved, it shall distinguish proposed, awaiting approval, initiated, and confirmed outcomes as applicable. It shall report an outcome as confirmed only when supported by the applicable execution evidence.

5.8.5.3 Where reliance on a Result depends on material assumptions, source limitations, unverified information, or an unavailable capability, the service shall communicate the relevant limitation before or with the Result in the Interface used for the decision. It shall distinguish information generated or inferred by the service from information verified against an identified source where that distinction materially affects reliance.

5.8.5.4 Where a change in mode, permissions, available context, or connected tools materially changes what the service can do or what its Result establishes, the service shall communicate the change before the user authorizes affected Work or receives a Result for reliance. Where Work cannot continue as requested, the service shall identify what remains possible and any decision or action required from the user.

5.8.5.5 The responsible party shall specify acceptance criteria for whether representative users can correctly identify the applicable capability, action status, limitations, and next step under intended conditions of use. It shall retain Evidence of the information presented, its timing, the Interface and configuration examined, and the basis for determining whether the criteria were met.

5.9 Factor 9 — Competence of Personnel

Indicator 9.1 — Identification of Object-Critical Roles

5.9.1.1 The responsible party shall identify in the Description, or in retrievable records referenced by it, personnel roles whose performance can materially affect realization of the Intended Functions or Intended Results, or the prevention or limitation of Harm. The identification shall cover roles performed by the provider and by external contributors and shall account for normal operation, material changes, and abnormal conditions.

5.9.1.2 As applicable, the identification shall cover personnel who can approve changes to models, instructions, data sources, or tool permissions, authorize consequential actions, review escalated Results, stop or take over Work, reconcile uncertain external outcomes, or approve resumption. A role shall be identified by the decisions and Work it performs, including when those responsibilities are shared or assigned to a supplier.

5.9.1.3 For each identified role, the Description shall specify the affected Work and Results, decisions reserved to the role, authority and Interfaces needed to act, and the conditions requiring availability or a qualified substitute. Where the service depends on the role to intervene before a specified point, the Description shall identify the latest effective point and the service response if the role is unavailable.

5.9.1.4 The responsible party shall review the identification when service autonomy, authority, staffing arrangements, or external contributions change materially, or when Evidence shows that an omitted or unclear assignment can affect the Work or Result.

Indicator 9.2 — Role-Appropriate Technical and Contextual Competence

5.9.2.1 For each role identified under 5.9.1, the responsible party shall specify in the Description, or in retrievable records referenced by it, the technical and contextual competence required for the assigned Work and decisions. The specification shall identify the service configurations, intended conditions of use, and abnormal conditions within the role's scope, together with the observable performance criteria and Evidence needed to establish competence.

5.9.2.2 Personnel who evaluate or approve changes to models, instructions, data sources, or tool permissions shall be able, within their assigned scope, to identify affected Intended Functions and assistance limits, interpret relevant evaluation Evidence, and recognize gaps that prevent a release decision. Where a change alters what the service can execute, the competence criteria shall address the ability to check that permissions and human approval controls still match the authorized scope.

5.9.2.3 Personnel who review escalated Results or intervene in ongoing Work shall be able, within their assigned scope, to recognize when a generated Result relies on an unresolved material assumption, distinguish proposed and initiated actions from confirmed outcomes, and use the applicable controls to restrict, stop, or take over Work within the required time. Where assigned to reconcile or resume Work, they shall be able to verify external action status and determine whether context, corrections, and authorization remain valid before resumption.

5.9.2.4 The contextual competence specification shall address the domain knowledge needed to judge consequences, the needs and circumstances of affected recipients, and the limits of the role's authority and expertise. It shall identify when the role must refer a decision to a specialist or other authorized person and what Work may proceed while that decision remains unresolved.

5.9.2.5 Where a critical decision depends on contributions from several roles or suppliers, the responsible party shall specify which competence each contributor must provide and which role shall integrate the findings and make the decision. The specification shall identify how material disagreements or gaps shall be resolved before the decision is made.

Indicator 9.3 — Competence for Abnormal, Degraded, and Non-Routine Conditions

5.9.3.1 For each role identified under 5.9.1 that must act under abnormal, degraded, or non-routine conditions, the responsible party shall specify in the Description or in retrievable records referenced by it the scenarios, decisions, available controls, and time limits for which competence shall be demonstrated. Scenario types may be described under 4.3.10. Material differences in evidence, authority, or available controls shall be identified.

5.9.3.2 Where personnel must intervene, they shall be able to use the evidence available at the time to distinguish a confirmed abnormal condition from an unconfirmed prerequisite under 5.8.4.1, identify affected Work, and apply the authorized restriction, stop, or handoff before the specified point. The criteria shall cover, as applicable, stale conversation context, an unavailable approval channel, and fallback configurations that lack required controls.

5.9.3.3 Where personnel must reconcile or resume Work following an interruption, they shall be able to verify the status of external actions, identify any action that could be duplicated, and establish whether current instructions, corrections, authorizations, and human controls permit resumption. Where required evidence or authority is unavailable, they shall be able to maintain the restriction and escalate to the specified role.

5.9.3.4 Before assigning unsupervised responsibility for these decisions, the responsible party shall establish competence using observed performance in representative controlled exercises or documented performance under comparable conditions. The Evidence shall identify the scenario, service configuration, role, actions taken, timing where material, and outcome against the specified criteria.

Indicator 9.4 — Recognition of Limits of Authority and Expertise

5.9.4.1 For each role identified under 5.9.1, the responsible party shall specify in the Description or in retrievable records referenced by it the limits of decision authority and demonstrated expertise, the conditions requiring referral, the role authorized to receive the referral, and the Work permitted while the decision is pending.

5.9.4.2 Personnel shall be able to recognize when an AI-generated recommendation, confidence indication, or tool response requires evidence or specialist judgment beyond their assigned expertise. They shall also be able to recognize when available system permissions exceed their decision authority, including permission to expand tool access, override a safeguard, or resume restricted Work.

5.9.4.3 Where a required decision exceeds either limit, personnel shall refer it with the affected Work, the decision needed, the available Evidence, and material uncertainties. The responsible party shall ensure that the decision-dependent Work remains within its specified limits until an appropriately authorized and competent person resolves the referral. If the decision cannot be obtained by the last effective point, the specified restriction, suspension, or handoff shall take effect.

5.9.4.4 The responsible party shall retain Evidence that personnel can recognize and act within both limits under representative conditions. Records of material referrals shall identify the reason, the role accepting responsibility, the disposition, and any authorization required before affected Work resumes.

Indicator 9.5 — Evolution of Competence with Lifecycle and Context

5.9.5.1 The responsible party shall reassess the competence required for affected roles before a planned transition in lifecycle stage, service configuration, or context of use that materially changes their Work or decisions. As applicable, the review shall address changes in model behavior, tool permissions, assistance limits, recipient needs, and the authority or time available for human intervention.

5.9.5.2 The responsible party shall compare the revised criteria under 5.9.2 and 5.9.3 with the Evidence for each affected assignment. Where earlier Evidence is used, the responsible party shall document its applicability to the changed duties and establish competence for any material differences before unsupervised Work is authorized.

5.9.5.3 Where an unplanned change, incident, or observed performance identifies a possible competence gap, the responsible party shall restrict the affected assignment pending review. Where a gap remains, the responsible party shall provide qualified supervision, reassign the Work, or suspend it until the required competence has been demonstrated.

5.9.5.4 The responsible party shall retain retrievable Evidence linking the change, affected roles, revised criteria, competence determinations, and resulting assignment decisions. The records shall identify when restrictions and authorizations took effect and the service configurations and conditions they cover.

Indicator 9.6 — Competence-Related Risk Awareness

5.9.6.1 For each role identified under 5.9.1, the responsible party shall specify the material pathways to Harm that can arise from a missed, incorrect, or delayed decision within the role's scope. The Description or referenced records shall identify the observable signs personnel can use, the Work and persons affected, the required response, and the latest effective point for intervention where material.

5.9.6.2 Personnel shall be able to explain and demonstrate the response required when the AI service presents plausible but unsupported information, conceals a material uncertainty through an apparently complete Result, or reports an external action whose outcome remains unconfirmed, as applicable to their role. They shall distinguish the service's report from Evidence sufficient to support reliance or further action.

5.9.6.3 Where personnel identify a material gap in their own or the assigned team's competence, they shall report it to the designated role and apply the specified limits on affected Work while the gap is addressed. Where a decision depends on specialist judgment that cannot be obtained in time, they shall apply the specified restriction, suspension, or handoff.

5.9.6.4 The responsible party shall verify role-specific risk awareness using representative decisions or exercises and retain Evidence of the signs recognized, the potential consequences identified, and the response taken against the applicable criteria. Findings that reveal a material gap shall be addressed under 5.9.5.

5.10 Factor 10 — Regulatory and Standards Compliance

Indicator 10.1 — Identification of Applicable Regulatory and Standards Obligations

5.10.1.1 The responsible party shall identify and maintain in the Description or in retrievable records referenced by it the regulatory and standards obligations applicable to the AI Service Quality Object. For each obligation, the record shall identify the source and applicable provision, effective date or edition, basis of applicability, and affected service functions, deployments, lifecycle stages, and roles.

5.10.1.2 The determination of applicability shall consider where and to whom the service is made available, the Intended Functions and uses, the data processed, the decisions or external actions the service supports or executes, and the role of each party in providing, deploying, or operating the service. Where these facts differ among deployments, the responsible party shall identify the material differences and their effect on applicability.

5.10.1.3 The record shall distinguish binding regulatory obligations, standards requirements made applicable by regulation or agreement, and requirements arising from voluntary commitments or claims. It shall identify any material unresolved question of applicability, who is responsible for resolving it, and the limits on affected Work or claims pending resolution.

5.10.1.4 The responsible party shall review the determination when an applicable source changes or when changes in deployment, Intended Function, autonomy, data use, or supplier responsibilities can affect applicability. The record shall identify the date, basis, and outcome of each material review.

Indicator 10.2 — Integration of Regulatory and Standards Requirements into Object Decisions

5.10.2.1 The responsible party shall translate each applicable requirement identified under Indicator 10.1 into decision criteria for the AI Service Work or Result it governs. The criteria shall identify the conditions that make the requirement applicable, the assistance or action permitted or required under those conditions, the point by which the decision shall take effect, and the party authorized to make or approve the decision.

5.10.2.2 Where a decision depends on facts established during an interaction, the responsible party shall specify which facts the AI Service may use, when it shall seek clarification or refer the decision to an authorized person, and what assistance may continue while applicability is unresolved. The AI Service shall withhold any consequential output, disclosure, or external

action whose permissibility depends on the unresolved facts until the applicable conditions are established or an authorized decision permits continuation.

5.10.2.3 Where a change in recipient circumstances, intended use, deployment context, or service configuration affects an applicable requirement, the responsible party shall reassess the affected decision criteria before the affected Work proceeds under the changed conditions. This shall include, as applicable, changes to the AI model, fallback configuration, connected tools, and external contributions. Where the applicable criteria cannot be met, the responsible party shall restrict or suspend the affected Work.

5.10.2.4 The responsible party shall maintain traceability from applicable requirements to the decision criteria and the service configurations and operating procedures that implement them. For consequential decisions governed by those criteria, the responsible party shall specify which facts, criteria versions, authorizations, and resulting actions or restrictions shall be recorded so that the basis for the decision can be reconstructed.

5.10.2.5 Where a recipient request involves information subject to an applicable access, use, or disclosure restriction, the AI Service shall establish that each restricted operation is permitted before performing it. If permission cannot be established or the operation is prohibited, the AI Service shall block or stop the affected search, retrieval, use, or disclosure, including through connected tools. The AI Service shall communicate the limitation without revealing protected information or confirming its existence where that confirmation is restricted. Assistance within the permitted scope may continue.

The responsible party shall document, in the Description or in retrievable records referenced by it, the categories of requests subject to applicable access, use, or disclosure restrictions, the basis and conditions for each restriction, and the required AI Service response.

Indicator 10.3 — Demonstration of Conformity with Applicable Requirements

5.10.3.1 The responsible party shall maintain or reference evidence sufficient to determine whether each requirement identified under Indicator 10.1 was met. The evidence shall identify the requirement and the AI Service Work or Result it covers, together with the relevant deployment, operating conditions, and period.

5.10.3.2 Where conformity depends on how the AI Service handled a recipient interaction, the evidence shall establish the sequence of events relevant to the requirement. As applicable, it shall establish the information available at the decision point, the service configuration in effect, the output actually delivered or the external action actually executed, and whether any required restriction, referral, or intervention took effect by the specified point.

5.10.3.3 The responsible party shall determine whether evidence remains applicable after a change to the model, service configuration, connected tools, or external contributions. Where a change can affect the service's ability to meet an applicable requirement, the responsible party shall obtain evidence covering the changed conditions before relying on a conformity determination for the affected Work.

5.10.3.4 The responsible party shall record whether each applicable requirement was met, was unmet, or could not be determined from the available evidence. Any conformity statement shall

be limited to the requirements, configurations, conditions, and period supported by the evidence.

Indicator 10.4 — Treatment of Conflicting or Overlapping Requirements

5.10.4.1 The responsible party shall identify applicable requirements that govern the same AI Service Work, Result, Interface, or operating condition and that overlap or may conflict. The responsible party shall document the source, jurisdiction where relevant, basis of applicability, applicable edition or effective date, and conditions of each requirement.

5.10.4.2 The responsible party shall evaluate how the requirements apply together to the affected Work. As applicable, the evaluation shall address requirements for retaining interaction Evidence and limiting retention, providing notice and restricting disclosure, and meeting delivery deadlines while obtaining required authorization or review.

5.10.4.3 The responsible party shall document how the proposed resolution fulfills the applicable obligations and affected Functional Requirements. Where clarification, approval, or an authorized alternative is required, the responsible party shall obtain it from the party with authority over the applicable obligation before relying on the resolution.

5.10.4.4 Where a conflict remains unresolved, the responsible party shall specify the affected configurations and Work, the restrictions that apply, and the conditions for lifting them. The AI Service shall withhold the affected step where proceeding would depend on the unresolved conflict. Relevant Findings shall identify the unresolved matter and its effect on the compliance conclusion.

Informative example. A service must retain Evidence that an applicant authorized submission, while a separate applicable requirement limits retention of personal information. The responsible party determines whether a bounded authorization record satisfies both requirements. If it does, the responsible party specifies what the record contains, who may retrieve it, and when it expires. If the obligations cannot be reconciled for that workflow, the affected submission step is restricted pending resolution.

Indicator 10.5 — Continued Validity of Compliance Under Change

5.10.5.1 Before relying on an existing approval, assessment, or compliance Finding for changed Work, the responsible party shall determine whether it remains applicable to the changed AI Service. The review shall address, as applicable, changes to the model or routing configuration, connected tools, data handling, deployment conditions, intended use, and external contributions.

5.10.5.2 Where the source of an applicable obligation changes, the responsible party shall determine its applicability to the AI Service, its effective date, and any transition conditions. Where the service changes provider, deployment jurisdiction, or operating arrangement, the responsible party shall determine whether any approval or authorized alternative remains applicable.

5.10.5.3 The responsible party shall specify when any required authorization, reassessment, or notification shall be completed. Where completion is a condition for proceeding, the affected Work shall remain restricted until the condition is met. The responsible party shall evaluate

temporary configurations, testing in operation, and emergency fallbacks against the requirements applicable to their actual use.

5.10.5.4 The responsible party shall update the applicable requirements, responsibility assignments, operating restrictions, and supporting records. Superseded approvals and Findings shall remain distinguishable from those applicable to the current configuration.

Informative example. An AI Service replaces its application submission tool. The prior compliance Finding relied on the tool preserving Evidence of authorization and linking the agency acknowledgment to the same submission. Before relying on that Finding for the replacement, the responsible party establishes whether the Evidence remains available and traceable. If the service cannot establish the required basis for submission, it may continue drafting while submission remains restricted.

Indicator 10.6 — Recognition of Compliance Limitations

5.10.6.1 A compliance statement shall identify or traceably reference the obligations addressed, the AI Service or part examined, the configurations and operating conditions covered, and the relevant period.

5.10.6.2 The statement shall distinguish demonstrated fulfillment, demonstrated nonfulfillment, requirements or conditions outside the examination scope, and matters for which the Evidence is insufficient to reach a conclusion.

5.10.6.3 Where the statement relies on a conditional approval, exemption, authorized alternative, or temporary permission, it shall preserve the applicable scope, conditions, limits, and duration. Where fulfillment depends on a connected tool or external party, the statement shall distinguish the assignment of responsibility from the Evidence of fulfillment.

5.10.6.4 The scope of a compliance conclusion shall be limited to the obligations and conditions supported by the Evidence. A broader Claim concerning Quality, Safety, or fulfillment of Intended Functions and Intended Results shall have its own supporting Evidence and Findings. Material limitations shall be reflected in the relevant Findings and any resulting Quality Claim.

Informative example. A compliance Finding covers authorization and submission records for the primary application tool during the examined period. The fallback tool was outside the examination scope. A statement based on that Finding identifies the primary tool and period and makes clear that the fallback tool was outside the scope.

5.11 Factor 11 — Object and Stakeholders

The AI Service Quality Object includes the recipient's participation in the realization of its Intended Functions and Intended Results. Other persons or organizations can be affected by the service's Work or Results even when they take no part in the interaction. The stakeholders relevant to a particular interaction can change as the recipient clarifies the request, supplies information about another person, or asks the service to take an external action.

Indicator 11.1 — Identification of Affected Stakeholders and Communities

5.11.1.1 The responsible party shall specify, for each Intended Function, the stakeholder roles relevant to its realization and effects. The specification shall distinguish the person interacting

with the service, the intended recipient of the Result, any person on whose behalf the request is made, any person authorized to approve or intervene in the Work, and other persons whose information, access, or interests can be affected. A person may occupy more than one role.

Informative example. A daughter asks the AI Service to help prepare her father's benefits application. She is the person interacting with the service and receives the draft. Her father is the applicant on whose behalf the Work is performed. The agency receives the submitted application. The father or his authorized representative may authorize submission. A household member whose income information is disclosed may be affected even without interacting with the service.

5.11.1.2 The responsible party shall specify the circumstances in which the AI Service shall establish or reassess a stakeholder's role during an interaction. These shall include, as applicable, a request made on behalf of another person, a change in the intended recipient of an output, and a transition from preparing an output to sending it or taking an external action.

5.11.1.3 Where the permitted scope of assistance depends on a stakeholder's role or authority, the AI Service shall establish that role or authority to the degree necessary for the affected Work before proceeding. The responsible party shall specify what assistance may continue while the role or authority remains unresolved.

Informative example. The service can identify the daughter and applicant roles in its first response and begin drafting. If the daughter later asks it to submit the application, the service establishes the required authority before invoking the submission tool. If another person must confirm, the service holds submission while awaiting confirmation. Drafting, waiting for authorization, and waiting for the agency's acknowledgment are distinct intervals.

5.11.1.4 The AI Service shall request personal identification only where the applicable Work requires it. Where the relevant stakeholder role can be established without identifying the person, the service shall use the role information sufficient for the Work.

5.11.1.5 For each Intended Function, the responsible party shall identify the persons, groups, and communities that can be materially affected by the AI Service's Work or Results, including those who do not interact with the service. Where individual identities are unnecessary or unavailable, the responsible party shall describe the affected groups by their relationship to the Work or by the conditions under which they may experience its effects.

Indicator 11.2 — Identification of Object Impacts Relevant to Stakeholders and Communities

5.11.2.1 For each Intended Function, the responsible party shall identify the material beneficial and adverse effects that can arise from the AI Service's Work or Results, the stakeholders who may experience those effects, and the conditions under which the effects can occur. The description shall address, as applicable, effects arising from generated outputs, decisions made in reliance on them, and actions performed through connected tools.

5.11.2.2 The responsible party shall distinguish observed effects from potential effects and document the material uncertainties that affect the identification of either. Where an effect depends on the stakeholder's circumstances or on action by another party, the description shall identify that dependency.

Informative example. An AI Service helps prepare a benefits application. An accessible draft can help the applicant review it independently. An incorrect statement about a household member can affect the application, and sending the draft to an unauthorized recipient can expose that person's information. Each effect concerns a different part of the Work and may affect a different stakeholder.

Indicator 11.3 — Alignment with Declared Commitments and Expectations

5.11.3.1 For each Intended Function, the responsible party shall specify categories of stakeholder needs and expectations that can materially affect the Intended Result, its delivery, or the permitted scope of assistance. The specification shall define the conditions that require clarification, confirmation, or referral, including how the AI Service shall respond when a need or expectation expressed during an interaction is outside the specified categories or its effect on the Work is unclear. The service shall resolve the uncertainty under 5.11.3.2 or 5.11.3.3, as applicable, before proceeding with the affected step. Other Work may continue where it can proceed independently within the permitted scope of assistance.

5.11.3.2 Where an unresolved stakeholder need or expectation could materially change the Result or its effects, the AI Service shall seek clarification before the first step that depends on it. The responsible party shall specify when clarification may be obtained from the person interacting with the service and when confirmation from the affected stakeholder or an authorized representative is required.

5.11.3.3 Where relevant stakeholder needs or expectations conflict, the AI Service shall apply the specified decision criteria for the affected Work. Where those criteria require a stakeholder or authorized person to resolve the conflict, the service shall communicate the unresolved issue to the appropriate person and limit its Work to what may proceed pending resolution.

5.11.3.4 Where a recipient corrects a stakeholder need or expectation during an interaction, the AI Service shall apply the correction to the affected Work still in progress. If an affected Result has already been delivered or an external action has already been executed, the service shall communicate its status and identify the corrective steps available within the service's scope.

5.11.3.5 The responsible party shall specify and document the review interval for the description required by 5.11.3.1, taking into account the consequences of incorrectly determining or handling a stakeholder need or expectation. The responsible party shall also review the affected provisions before implementing a planned change that can alter stakeholder needs, their effects, or the service's ability to handle them, and when an interaction, complaint, or incident indicates that the description may be inadequate. Where the review identifies a gap, the responsible party shall update the affected provisions and determine whether any Work shall be restricted pending implementation of the update.

Informative example. A daughter asks the AI Service to prepare her father's benefits application and send updates to her email. Her request establishes her preferred contact channel. If the service specification requires the applicant or an authorized representative to confirm where notices should go, the service may continue preparing the draft while holding the contact designation for confirmation before submission.

5.11.3.6 For each Intended Function, the responsible party shall specify how established stakeholder needs and expectations affect the AI Service's choice of assistance, the content and delivery of the Result, and any external action. The specification shall distinguish mandatory conditions from preferences that may be adjusted and identify who may authorize an adjustment.

5.11.3.7 Where the AI Service relies on a stakeholder instruction from an earlier interaction, the service shall apply the specified conditions for reuse. Before a consequential action, the service shall reconfirm the instruction where its applicability to the current Work cannot be established. The responsible party shall specify the circumstances requiring reconfirmation, including, as applicable, a change in the intended recipient, the purpose of the action, or the scope of authorization.

5.11.3.8 Where the AI Service delegates Work to a connected tool or another service, it shall carry forward the stakeholder conditions applicable to that Work. If the receiving tool or service cannot honor a mandatory condition, the AI Service shall withhold the affected action and communicate the limitation before offering any available option within the permitted scope of assistance.

Informative example. A daughter asks the AI Service to prepare her father's benefits application. The father authorizes the service to submit the application to the agency and asks to receive a copy. When the daughter resumes the conversation and asks the service to send a copy to a new caregiver, the service establishes whether that separate disclosure is authorized before sending it. If the connected tool cannot deliver the father's copy as specified, the service handles that limitation according to whether delivery of the copy is a condition of submission or a separate preference.

5.11.3.9 The responsible party shall identify the declared commitments relevant to each Intended Function and specify how they govern the affected Work and Results. Where a material stakeholder expectation conflicts with a declared commitment or an applicable requirement, the AI Service shall communicate the resulting limitation before the affected Work proceeds. Any adjustment shall be authorized within the applicable requirements and the permitted scope of assistance.

Indicator 11.4 — Consideration of Cumulative and Long-Term Community Effects

5.11.4.1 The responsible party shall identify material effects that may arise through repeated interactions, repeated use of AI-generated Results, or the accumulation of external actions over time. The description shall identify the affected stakeholders, the Work or Results that contribute to the effect, and the period or operating conditions relevant to its development.

5.11.4.2 For each identified cumulative or long-term effect requiring observation, the responsible party shall specify the information needed to recognize it, the period over which it shall be examined, and the conditions that trigger review or a change in the affected Work. Where the service cannot observe a material part of the effect, the responsible party shall document that limitation and specify what conclusions the available information can support.

Informative example. An AI Service repeatedly directs applicants whose information does not fit standard form fields to an alternative process. Each referral may be appropriate in isolation.

Over time, repeated referrals with no successful handoff may delay access for a group of applicants. Examination therefore considers referral outcomes over a defined period as well as individual responses.

Indicator 11.5 — Responsiveness to Object-Related Community Concerns

5.11.5.1 The responsible party shall specify how recipients and other affected stakeholders can raise concerns about the AI Service's Work, Results, or effects. The specification shall identify the information needed to locate the affected Work, the party responsible for review, and the response times or escalation conditions appropriate to the consequences.

5.11.5.2 Where a concern indicates that an external action may be incorrect, unauthorized, or otherwise inconsistent with applicable requirements, the AI Service shall apply the specified restrictions to pending dependent Work. Where the action has already taken effect, the service shall communicate its known status and the available route for review or correction within applicable disclosure limits.

5.11.5.3 The responsible party shall record the disposition of material concerns and determine whether recurring or related concerns require review of the affected requirements, configurations, or monitoring arrangements.

Indicator 11.6 — Transparency of Object Impacts and Limitations

5.11.6.1 The responsible party shall specify which material effects, uncertainties, and limitations shall be communicated to each affected stakeholder role, in what form, and by what point. Where the information could change a stakeholder's decision to authorize, use, or rely on a Result, the AI Service shall communicate it before that decision where the service controls the timing.

5.11.6.2 Where the AI Service cannot observe an external outcome, reach an affected stakeholder through an authorized channel, or correct a Result that has already taken effect, it shall communicate the applicable limitation and any available next step within the permitted scope of assistance and applicable disclosure limits.

5.11.6.3 For each Intended Function, the responsible party shall specify which stakeholder roles require communication, the events that trigger it, the information to be conveyed, the authorized recipient and channel, and the point by which the communication is required. The specification shall distinguish communication that the AI Service can deliver from communication that depends on an authorized person or an external service.

5.11.6.4 Where a stakeholder must confirm, correct, or intervene before the Work proceeds, the AI Service shall communicate the proposed action, the decision required, and the status of the affected Work in a form the stakeholder can use. Where the service cannot reach the stakeholder through an authorized channel, it shall communicate the limitation to the person interacting with the service, within the applicable disclosure limits, and hold the dependent step.

5.11.6.5 Where a communication is part of the Intended Result, the AI Service shall report its status using the evidence available. The reported status shall distinguish preparation, dispatch, confirmed delivery, and any acknowledgment required for completion. If delivery fails or cannot

be established, the service shall communicate that status and identify any next step available within its scope.

Informative example. After submitting a benefits application, the AI Service receives an agency receipt. The applicant has also asked for a copy. The service reports that the agency received the application and separately reports whether the applicant's copy was sent or delivered. If the applicant's authorized delivery channel fails, the service identifies the failure and the available way to provide the copy without disclosing it to another person.

5.12 Factor 12 — Object Monitoring

Monitoring examines the AI Service Quality Object during operation. Its scope includes the Work performed, the Results delivered, and the effects that can be observed. Depending on the Intended Function, a departure may be visible in a single response, in a sequence of interactions, or in the outcome of an action performed through a connected tool.

Indicator 12.1 — Identification of Monitoring-Relevant Object Behaviors

5.12.1.1 For each Intended Function, the responsible party shall specify the departures from applicable requirements or Quality Outcome Criteria that monitoring is intended to detect. As applicable, the specification shall cover generated outputs, changes in the service's handling of a request across interactions, actions performed through connected tools, and discrepancies between the status reported to the recipient and the status recorded by the connected system.

5.12.1.2 For each departure to be detected, the responsible party shall specify the observable events or conditions, the sources of monitoring information, the scope and frequency of observation, and the party responsible for reviewing the results. Where a departure can be detected only by relating events across interactions or systems, the specification shall identify which events must be associated and what evidence establishes that they concern the same Work.

5.12.1.3 The responsible party shall specify the latest point at which a departure must be detected to permit the required response. Where a response is required before an output is delivered or an external action takes effect, the monitoring and response arrangements shall operate before that point. Where the purpose is to detect a recurring pattern, the specification shall define the observation period and the conditions that trigger review.

5.12.1.4 Where monitoring depends on information from a connected tool or external service, the responsible party shall specify how missing, delayed, or conflicting information is identified and how it affects the monitoring conclusion. Missing confirmation shall be distinguished from evidence that the monitored action succeeded or failed.

5.12.1.5 For each monitored behavior or condition, the responsible party shall identify the expected behavior or applicable limit, the part of the Work at which it can be observed, and how material departures shall reach the party responsible for responding. Where relevant to Function Realization or protective effectiveness, the specification shall include routing and fallback modes, availability or bypass of protective controls, and loss or delay of status information from connected tools.

Informative example. An AI Service submits a benefits application through a connected tool. The monitoring specification identifies three events to associate with the same submission: the submission request, the agency acknowledgment, and the status shown to the recipient. If the acknowledgment is missing, the monitoring record distinguishes an unconfirmed submission from a rejected submission. If the service reports receipt before the agency acknowledges it, the discrepancy is detectable from the event sequence.

Indicator 12.2 — Adequacy of Monitoring to Detect Degradation and Deviation

5.12.2.1 The responsible party shall carry out the monitoring specified under Indicator 12.1 for the AI Service Work and Results within its defined scope. Where monitoring uses sampling, the responsible party shall specify how cases are selected, which operating conditions the sample is intended to cover, and which conclusions the sample can support. Where a protective response is required before an output is delivered or an action takes effect, the required monitoring shall operate for each affected instance of Work.

5.12.2.2 The responsible party shall maintain the specified monitoring coverage when the AI Service changes its model, uses a fallback configuration, or transfers Work to a connected tool. Where a change prevents the service from performing monitoring required before a consequential step, the service shall withhold that step until the required monitoring is restored or an authorized alternative provides the required coverage.

5.12.2.3 The responsible party shall identify and record interruptions or gaps in monitoring, including the affected Work and period. Monitoring records and summaries shall distinguish Work for which monitoring detected no departure from Work for which the required observation was unavailable or incomplete.

5.12.2.4 The responsible party shall establish that monitoring can detect material departures within the time available for the required response. The basis shall account for the accuracy, sensitivity, and frequency of observation and the effect of thresholds, filtering, or delays. Where a departure can be confined to a configuration, interaction path, or group of recipients, the observation shall be capable of revealing it at that level.

5.12.2.5 Where a failure can affect both the Work and the monitoring relied upon to observe it, the responsible party shall establish that the required monitoring remains effective or specify an alternative observation or restriction. The service shall distinguish a requested or inferred action from a confirmed action where that distinction affects a decision or required response.

Informative example. An AI Service switches to a fallback tool to submit benefits applications. The fallback tool can transmit the application but cannot return the evidence needed to associate the agency acknowledgment with that submission. If the missing observation is required before submission may proceed, the service holds the submission. If it is required only to confirm receipt afterward, the service may submit under the specified conditions and report receipt as unconfirmed until supporting evidence is available.

Indicator 12.3 — Interpretation of Monitoring Information in Context

5.12.3.1 For each departure or monitoring condition requiring a response under Indicator 12.1, the responsible party shall specify the action to be taken, the party authorized to take it, and the

point by which it shall take effect. The specification shall distinguish actions the AI Service may take during the Work from cases requiring review or intervention by an authorized person.

5.12.3.2 When monitoring indicates a departure or leaves the outcome of an external action uncertain, the AI Service shall apply the specified restrictions to any dependent Work. Where retrying could duplicate or conflict with an action already performed, the service shall establish the status of the original action or use a specified retry procedure that prevents duplicate or conflicting effects. While the outcome remains uncertain, the service shall communicate that status to the recipient within the applicable disclosure limits.

5.12.3.3 Where monitoring identifies a recurring departure or a change in its frequency or consequences that meets the specified review criteria, the responsible party shall determine which Intended Functions, operating configurations, and Results may be affected. The responsible party shall document the response, any interim restrictions, and the evidence required to determine whether the response is effective.

5.12.3.4 The responsible party shall interpret monitoring information against the applicable requirements and the configuration, operating mode, and external contributions in effect when it was obtained. The interpretation shall distinguish expected variability, planned transitions, degradation, and faults in the monitoring or reporting path. Material inconsistencies shall be examined before relying on an interpretation that could affect consequential Work.

5.12.3.5 Monitoring Findings shall be traceable to the relevant observations, event sequence, and reasoning. Where a Finding depends on a derived measure, the method, inputs, and material assumptions or uncertainties shall be identifiable. A cleared alert or a return to the expected range shall establish resolution only where the available Evidence also supports that conclusion about the underlying condition.

Informative example. The AI Service sends a benefits application through a connected tool, but the acknowledgment times out. The service reports that receipt is unconfirmed and checks the status of the original submission before sending again, unless the specified retry procedure prevents duplicate submissions. If timeouts recur after a tool change and meet the review criterion, the responsible party determines which submissions may be affected and whether to restrict use of that tool while the issue is examined.

Indicator 12.4 — Detection of Emerging and Systemic Risk

5.12.4.1 The responsible party shall examine monitoring information for trends and combinations of events that could indicate developing loss of AI Service capability or protective effectiveness. As applicable, the examination shall address increasing fallback use, repeated loss of conversation context, recurring requests for human intervention, repeated bypasses of protective controls, and growing discrepancies between actions reported to recipients and outcomes confirmed by connected systems.

5.12.4.2 The responsible party shall examine related events across interactions, configurations, or connected tools for shared causes or dependencies. The examination shall consider whether events closed individually form a material pattern when viewed together, including where a shared model, data source, routing rule, or external service could affect apparently separate workflows.

5.12.4.3 The responsible party shall specify which warning patterns require further examination or precautionary action and when those patterns shall reach the party authorized to respond. The reported pattern shall identify the affected Work, the Evidence supporting its significance, and any material uncertainty about the cause or extent. Where the Evidence cannot distinguish between causes that would require different responses, the responsible party shall identify the investigation or interim restriction needed.

Informative example. A benefits application service resolves several individual acknowledgment timeouts. Monitoring across workflows shows that the timeouts are becoming more frequent in services using the same connected gateway. The responsible party examines the shared dependency, identifies which submissions may be affected, and determines whether to restrict retries while investigating.

Indicator 12.5 — Monitoring Relevance Under Change

5.12.5.1 The responsible party shall evaluate changes for their effects on monitoring coverage, expected behavior or limits, detection capability, interpretation, and response arrangements. As applicable, the evaluation shall address changes to models, routing, tool Interfaces, operating modes, Intended Functions, and the populations or conditions of use covered by monitoring.

5.12.5.2 Before relying on changed monitoring, the responsible party shall establish that it identifies the relevant events, associates them with the affected Work, and reaches the authorized party or protective response in time. Where a transition creates a monitoring gap, the responsible party shall specify its scope and duration, any alternative observation or operating restriction, and the conditions for restoring normal operation.

5.12.5.3 Monitoring records shall preserve the relationship between observations and the configuration, settings, and methods applicable when the observations were obtained. Where a change affects the meaning or comparability of a measure, the responsible party shall identify that difference in any comparison or trend Finding.

Informative example. A replacement submission tool returns an immediate message when it queues an application. The previous tool returned a message only after agency acknowledgment. Before relying on monitoring for the replacement, the responsible party establishes that it distinguishes queue acceptance from agency receipt and that any historical comparison uses equivalent events or explains the difference.

5.13 Factor 13 — Documentation and Traceability

Indicator 13.1 — Documentation of Object-Relevant Decisions and Assumptions

5.13.1.1 The responsible party shall document decisions and assumptions that materially affect the AI Service’s Intended Functions, Intended Results, applicable requirements, or permitted scope of assistance. As applicable, these shall include decisions concerning recipient participation, model selection and routing, fallback behavior, reliance on connected tools or external parties, and the points at which the service shall seek clarification, obtain authorization, or withhold an action.

5.13.1.2 The record shall identify the decision, its rationale and supporting basis, the party responsible for it, and the configurations and conditions to which it applies. For a material assumption, the record shall identify what was assumed, which Work or Result depends on it, and the conditions that require it to be reconsidered.

5.13.1.3 Where the AI Service makes a consequential decision during an interaction, the responsible party shall specify which facts, applicable criteria, and unresolved assumptions shall be recorded to establish the basis of the decision. The record shall distinguish information confirmed by an authorized source from information inferred by the service or supplied without confirmation where that distinction affects the permitted Work.

5.13.1.4 Where an assumption ceases to hold or a decision is superseded, the responsible party shall document the change and identify the affected Work, configurations, and applicable restrictions. The records shall allow an authorized examiner to determine which decision and assumptions applied when the Work was performed.

Informative example. A benefits application service is configured to report agency receipt only after it receives an acknowledgment linked to the submission. The decision record identifies the assumption that the connected tool can distinguish an agency acknowledgment from its own queue confirmation. If a replacement tool reports only queue acceptance, the responsible party documents the changed assumption and limits the status the service may report.

Indicator 13.2 — Traceability Between Functional Intent, Technical Requirements, and Realized Object

5.13.2.1 The responsible party shall maintain traceable links from each Intended Function and Intended Result to the applicable Functional and Technical Requirements and to the provisions through which they are realized. As applicable, those provisions shall include AI Service Constituents, Work, Interfaces, and necessary external contributions.

5.13.2.2 Where an Intended Function can be realized through more than one configuration or operating path, the traceability shall identify which requirements apply to each path and how they are fulfilled. This shall include, as applicable, primary and fallback models, connected tools, and handoffs to authorized personnel or external services.

5.13.2.3 For consequential Work, the responsible party shall maintain records sufficient to trace a delivered Result or external action back to the applicable requirements and the configuration and operating path used at the relevant time. Where a connected system performs part of the Work, the records shall distinguish the request sent by the AI Service from the action or outcome confirmed by that system.

5.13.2.4 When a requirement or a provision for its realization changes, the responsible party shall update the affected links and identify the Intended Functions, Intended Results, and dependent Work affected by the change. The links applicable to earlier Work shall remain distinguishable from those applicable to the changed service.

Informative example. The Intended Result is a benefits application submitted with the applicant’s authorization. The traceable links connect that Result to the authorization and submission requirements, the authorization check, and the connected submission tool. For a particular case, the record identifies which tool was used and whether the agency confirmed

receipt. The same links allow the responsible party to identify which requirements and workflows must be reviewed if the tool is replaced.

Indicator 13.3 — Documentation of Uncertainty, Limitations, and Known Vulnerabilities

5.13.3.1 The responsible party shall document material uncertainties, known limitations, and known vulnerabilities that can affect the AI Service's Intended Functions, Intended Results, or protective controls. Each record shall identify the affected Work, the configurations and conditions in which the issue applies, its possible effects, and the Evidence supporting what is known.

5.13.3.2 The record shall distinguish a demonstrated limitation or vulnerability from an unresolved uncertainty. Where the extent of an issue remains uncertain, the responsible party shall identify the conditions examined, the conditions that remain unexamined, and any restrictions or additional checks required before the affected Work may proceed.

5.13.3.3 The responsible party shall identify who needs the documented information to operate, supervise, use, or examine the AI Service. Where a limitation or uncertainty could materially affect a recipient's decision to authorize, use, or rely on a Result, the service shall communicate the relevant information before that decision, where it controls the timing. Access to sensitive vulnerability details shall be limited to authorized parties.

5.13.3.4 The responsible party shall update the record when new Evidence changes the known scope, consequences, or status of the issue. Where the responsible party changes or lifts a restriction, the record shall identify the supporting Evidence, the configurations covered by the decision, and when the decision took effect.

Informative example. A benefits application service can retrieve the applicant's earlier authorization in its primary configuration. Testing shows that a fallback configuration sometimes resumes the conversation without that record. The responsible party documents the conditions in which the loss was observed and the configurations whose behavior remains uncertain. If the required authorization cannot be established, the service may continue drafting but shall withhold submission and tell the recipient what confirmation is needed.

Indicator 13.4 — Traceability of Assessments, Evaluations, and Conformity Statements

5.13.4.1 The responsible party shall maintain traceable records for assessments and evaluations used to determine AI Service Quality or conformity. The records shall identify the AI Service or part examined, the applicable requirements and criteria, the methods used, and the configurations, operating conditions, and period covered.

5.13.4.2 Each Finding shall be traceable to the Evidence and reasoning that support it. Where Evidence was obtained from sampled interactions, a test environment, or an external party, the records shall identify the source and any limits on applying that Evidence to the AI Service as operated.

5.13.4.3 Each conformity statement shall be traceable to the Findings that support it. The traceable record shall identify the scope and conditions of the statement and any exclusions, unresolved Findings, or limitations that affect its use.

5.13.4.4 Where the AI Service changes in a way that may affect an existing Finding or conformity statement, the responsible party shall use the traceable records to identify the conclusions requiring review. Where a conclusion is revised, restricted, or superseded, the responsible party shall record the disposition and preserve the link to the Evidence and service configuration applicable to each version.

Informative example. An assessment confirms that the AI Service checks the applicant's authorization before submitting a benefits application through its primary tool. The conformity statement links to the applicable requirement, the sampled authorization and submission records, and the tool configuration examined. When a fallback tool is introduced, the traceable record shows which Finding requires review and whether the existing statement still applies to the primary path.

Indicator 13.5 — Documentation of Changes and Their Rationale

5.13.5.1 The responsible party shall document changes that materially affect the AI Service's Intended Functions, Intended Results, applicable requirements, or protective controls. As applicable, this shall include changes to models, instructions governing service behavior, routing and fallback rules, connected tools and their permissions, and necessary external contributions.

5.13.5.2 The change record shall identify the previous and revised provisions, the reason for the change, the Evidence and material assumptions supporting the decision, the party that authorized it, and the conditions and time at which it took effect. The record shall identify the affected Work, including Work already in progress where the change alters how it may continue.

5.13.5.3 For a staged rollout, temporary configuration, emergency change, or rollback, the responsible party shall document which Work was subject to each configuration and during what period. Where an external party makes a material change, the responsible party shall document what is known about the change, any uncertainty about its scope or timing, and the resulting decision about continued use.

Informative example. A benefits application service introduces a replacement submission tool in stages. The change record explains why the tool was replaced, identifies when each group began using it, and specifies how pending submissions were handled. If the service rolls back to the previous tool, the record identifies when the rollback took effect and which submissions still require confirmation from the replacement tool.

Indicator 13.6 — Accessibility and Continuity of Quality-Critical Information

5.13.6.1 The responsible party shall identify the information necessary to perform, supervise, resume, or examine consequential Work. For each category of information, the responsible party shall specify who or what may retrieve it, the point by which it shall be available, and the period during which it shall remain retrievable. As applicable, this shall include current operating instructions, applicable requirements and restrictions, authorization records, the status of pending external actions, and Evidence supporting Findings.

5.13.6.2 The responsible party shall preserve access to the applicable version of quality-critical information through changes in models, tools, storage arrangements, and responsible parties. Retrieved information shall remain usable and attributable to the Work, configuration, and

period to which it applies. Where current and superseded versions coexist, the service shall distinguish which version governs the Work in progress.

5.13.6.3 Where the AI Service cannot retrieve or establish the applicability of information required for a consequential step, it shall withhold the dependent step and communicate the limitation within the applicable disclosure limits. The responsible party shall specify the conditions for resuming the Work, including, as applicable, recovery of the record, confirmation through an authorized source, or transfer to an authorized person.

Informative example. A benefits application service migrates its records while some submissions are awaiting agency acknowledgment. The replacement system must be able to retrieve the authorization and submission status for each pending case. If the service cannot establish whether an application was already sent, it holds any repeat submission until the original status can be confirmed. Authorized examiners can still retrieve the earlier records under the specified access and retention conditions.

5.14 Factor 14 — Object Adaptability

Indicator 14.1 — Capability of Object to Accommodate Required Changes

5.14.1.1 The responsible party shall identify the changes the AI Service must be capable of accommodating to continue realizing its Intended Functions and Intended Results. For each required change, the responsible party shall specify the conditions that trigger it, the time available to make it, and the requirements that shall continue to be fulfilled during and after the transition.

5.14.1.2 The responsible party shall determine whether the integrated AI Service can accommodate each required change under the specified conditions. As applicable, this determination shall cover model replacement, changes to connected tools or external contributions, changes in demand, and changes in the conditions under which recipients use the service. It shall account for the capabilities and limitations of the Constituents, Work, and Interfaces on which the transition depends.

5.14.1.3 Where a change affects Work in progress, the responsible party shall specify which Work may continue, which shall be held or transferred, and the conditions for resumption. The AI Service shall preserve applicable authorization, protective controls, and the known status of pending external actions. Where the status of an action cannot be established, the service shall restrict dependent Work as required under Indicator 12.3.

5.14.1.4 Before relying on a changed operating path, the responsible party shall obtain Evidence that it can fulfill the applicable requirements under the specified conditions, including any conditions specific to the transition. Where capability has been demonstrated only within defined limits, the responsible party shall specify those limits and restrict operation accordingly.

Informative example. A benefits application service must replace a submission tool before its provider withdraws it. The service can continue preparing drafts during the transition. Before opening the replacement submission path, the responsible party establishes that it preserves the required authorization checks and can distinguish applications queued for delivery from

those acknowledged by the agency. Pending submissions remain linked to the original tool until their status is established.

6. Examination Basis

6.1 Identification of the AI Service Quality Object

6.1.1 Before determining the Applicable Basis under 6.2 and the Examination Scope and Conditions under 6.3, the party responsible for the examination shall identify the AI Service Quality Object in a documented record. The record shall reference the applicable Quality Object Description and identify the service provider, deployment, and period to which the identification applies.

6.1.2 The identification shall state the Intended Functions and Intended Results of the object and identify the Boundaries, Interfaces, Constituents, and essential Work and Results relevant to their realization. It shall identify the role of recipients and the external contributions on which the service depends, including Work performed by connected tools, external services, or authorized personnel.

6.1.3 Where the service operates through materially different configurations or paths, the identification shall distinguish them to the extent necessary to establish which provisions governed the Work. As applicable, it shall identify model and tool versions, routing and fallback paths, assistance limits, authorization arrangements, and the periods in which each configuration applied.

6.1.4 Where the examination concerns a defined part of the service, the record shall identify that part, its relationship to the wider AI Service Quality Object, and the dependencies that remain relevant to its examination. Any material uncertainty about the identity, configuration, or Boundaries of the object shall be recorded and carried forward when establishing the Examination Scope and Conditions under 6.3.

Informative example. An examination concerns the submission function of a benefits application service during a period when its submission tool was replaced. The identification distinguishes the drafting and submission functions, identifies the configurations used before and after the change, and includes the agency acknowledgment Interface because it is necessary to establish receipt. Clause 6.3 then specifies which operating paths and conditions the examination covers.

6.2 Applicable Basis

6.2.1 Before conducting the examination, the party responsible for the examination shall establish and document the Applicable Basis for the AI Service Quality Object identified under 6.1. The Applicable Basis shall identify the requirements and criteria against which the object or the examined part will be evaluated.

6.2.2 The Applicable Basis shall include, as relevant, the requirements of this Standard, the Functional and Technical Requirements for the identified object, and requirements arising from applicable regulatory obligations, agreements, or declared commitments. The record shall

identify the source and applicable version of each requirement, why it applies, and the Intended Function, Intended Result, Work, or Interface to which it relates.

6.2.3 Where a requirement needs an examination criterion, the party responsible for the examination shall establish that criterion before evaluating the Evidence against it. The criterion shall specify the observable condition or result and any applicable limit, tolerance, or decision rule. Where different criteria apply to primary, fallback, or externally performed Work, the Applicable Basis shall state which criterion applies to each path and under what conditions.

6.2.4 The record shall distinguish requirements that do not apply to the identified object from applicable requirements that fall outside the Examination Scope established under 6.3. The party responsible for the examination shall document the reason for each determination. Where applicability remains unresolved, the record shall identify the issue and any resulting limitation on the Findings or Quality Claims.

6.2.5 Where a requirement or criterion changes during the period covered by the examination, the Applicable Basis shall identify which version applies to the affected Work and when it took effect. Where two applicable sources impose different conditions on the same Work, the record shall identify both and document how the difference was resolved or what remains unresolved.

Informative example. A benefits application service may prepare a draft before the applicant authorizes submission. The Applicable Basis specifies that valid authorization must exist before the service sends the submission request. For the separate requirement to report agency receipt, it specifies that the service may report receipt only after receiving an agency acknowledgment linked to that submission. If a replacement tool confirms only that it queued the request, that confirmation supports a queued status. Clause 6.3 then establishes whether the examination covers the primary tool, the replacement tool, or both.

6.3 Examination Scope and Conditions

6.3.1 The party responsible for the examination shall establish and document its scope and conditions before evaluating the Evidence. The scope shall identify the AI Service Quality Object or part examined under 6.1, the requirements from the Applicable Basis under 6.2 that will be examined, and the period covered.

6.3.2 The scope shall identify the Intended Functions and Intended Results, Work, Constituents, Interfaces, and external contributions covered by the examination. Where the examination covers only part of the AI Service Quality Object, it shall include the dependencies necessary to evaluate that part against the requirements in scope.

6.3.3 The documented conditions shall identify the service configurations, operating paths, and circumstances under which conformity will be examined. Where material to the requirements in scope, these shall include differences in model or tool versions, assistance limits, authorization arrangements, and conditions such as fallback, interruption, handoff, or changes during the period examined.

6.3.4 The party responsible for the examination shall document applicable requirements, configurations, operating paths, or conditions excluded from the scope and the reasons for their exclusion. Where Evidence for a matter within scope is unavailable or insufficient, the

examination record shall identify the resulting limitation on the Findings and any Quality Claims based on them.

6.3.5 If the scope or conditions change during the examination, the party responsible for the examination shall record the change, its reason, and its effect on Evidence already obtained and Findings already made. The final examination record shall identify the scope and conditions actually covered. Any resulting Quality Claim shall remain within the scope and conditions supported by the Findings.

Informative example. An examination of a benefits application service covers both the original submission tool and its replacement, including cases pending when the change took effect. The examination specifies how it will cover authorization before submission, confirmation that a request was queued, and acknowledgment that the agency received it. If agency acknowledgment records are unavailable for submissions made through the replacement tool, agency receipt remains within scope. The Findings record the Evidence gap, and any resulting Quality Claim is limited to the submission status the Evidence establishes.

7. Evidence

7.1 The party responsible for the examination shall identify the Evidence needed to evaluate the requirements and criteria established under 6.2 within the scope and conditions established under 6.3. The examination record shall link the Evidence obtained to the requirement or criterion it is used to evaluate and to the AI Service Work or Result it concerns.

7.2 For each item or set of Evidence relied upon, the examination record shall identify its source, how and when it was obtained, and the service configuration, operating path, and conditions it represents. Where Evidence is derived from original records by filtering, aggregation, redaction, or summarization, the record shall document any transformation material to its interpretation and retain a traceable link to the source.

7.3 The party responsible for the examination shall evaluate whether the Evidence is relevant, reliable, and sufficient for the requirement or criterion being examined. Where service behavior can vary across model runs, recipient inputs, or operating paths, the examination record shall identify how cases were selected and the limits of their coverage. Evidence from tests, simulations, or a different configuration shall be relied upon only to the extent that its applicability to the Work and conditions examined has been established.

7.4 Where an Intended Result depends on a connected tool, external service, or authorized person, the Evidence shall distinguish, as applicable, the proposed action, its authorization, its submission, its execution, and achievement of the Intended Result. Where the AI Service reports the status of externally performed Work, the Evidence shall allow the reported status to be compared with what the external source confirmed at the time of reporting.

7.5 The party responsible for the examination shall record missing, conflicting, or unverifiable Evidence that could affect a Finding. The record shall identify the requirement or criterion affected, any steps taken to resolve the issue, and the uncertainty that remains for consideration under Clause 8.

7.6 Evidence relied upon shall remain retrievable by authorized parties for the applicable retention period. Where access, disclosure, or retention restrictions limit independent examination of the Evidence, the examination record shall identify the restriction and its effect on what can be verified.

Informative example. A benefits application service records the applicant's authorization, the request sent to the submission tool, and the status displayed to the applicant. The replacement tool confirms that it queued the request. The agency later issues a separate acknowledgment. If the agency acknowledgment cannot be linked to the specific application, the tool record can support a Finding that the request was queued. The examination record carries the unresolved link forward for the Finding on agency receipt.

8. Findings and Synthesis

8.1 The party responsible for the examination shall derive each Finding by evaluating the Evidence under Clause 7 against the applicable requirement or criterion established under 6.2. The examination record shall identify the requirement or criterion, the Work or Result examined, the supporting Evidence, and the reasoning leading to the Finding.

8.2 Each Finding shall state whether the Evidence establishes that the requirement or criterion was fulfilled, establishes that it was not fulfilled, or leaves fulfillment unresolved. The Finding shall identify the configurations, operating paths, conditions, and period to which it applies. Where Evidence supports different determinations under different conditions, the Finding shall distinguish those conditions and their corresponding determinations.

8.3 Where missing, conflicting, or unverifiable Evidence could affect a Finding, the party responsible for the examination shall record how the issue was evaluated and what remains unresolved. Where a Finding applies to selected cases or repeated model runs, the record shall state the extent to which the result can be applied to other Work within the Examination Scope.

8.4 Where a requirement concerns how the AI Service performs or whether it achieves an Intended Result, the Finding shall state what the Evidence shows about actual performance. If the Evidence shows only that a control was documented or configured, the Finding shall state that fulfillment of the performance requirement remains unresolved.

8.5 The party responsible for the examination shall synthesize the relevant Findings to determine what they collectively establish about the realization of each Intended Function and Intended Result examined. The synthesis shall account for dependencies and interactions among the applicable requirements, Work, Constituents, Interfaces, and external contributions. Where Findings concern separate contributions to a Result, the synthesis shall identify the Findings that establish whether those contributions worked together to achieve that Result.

8.6 The synthesis shall preserve the effect of material nonconformities, unresolved Findings, and differences among operating paths or conditions. Where an unresolved dependency or a nonconformity limits a conclusion about an Intended Function or Intended Result, the synthesis shall identify the affected conclusion and the extent of the limitation.

8.7 The examination record shall link each synthesized conclusion to the Findings that support or qualify it. The synthesis shall state what is established, under which conditions, and with what limitations, so that the basis and permissible scope of any Quality Claim under Clause 9 can be determined.

Informative example. An examination of a benefits application service establishes that the replacement submission path checked authorization before sending requests and accurately reported that the requests were queued. Agency acknowledgments could be linked to applications sent through the original path, while receipt remained unresolved for applications sent through the replacement path. The synthesis therefore distinguishes verified agency receipt for the examined cases on the original path from verified queueing on the replacement path.

9. Quality Claims

9.1 The party issuing a Quality Claim shall identify the intended recipients and the decision or use the Claim is intended to support. The Claim shall address the aspects of the AI Service Quality Object relevant to that decision or use, within the limits of the supporting Findings and synthesis under Clause 8.

9.2 The Claim shall identify the issuing party, the date of issue, and the AI Service Quality Object or part to which it applies. It shall state what was established about the fulfillment of specified requirements or the realization of specified Intended Functions or Intended Results. It shall identify the applicable configurations, operating paths, conditions, and period to the extent necessary for recipients to determine whether the Claim applies to their intended use.

9.3 The scope and certainty of the Claim shall remain within what the supporting Findings and synthesis establish. The Claim shall disclose material exclusions, nonconformities, unresolved Findings, and limitations that affect reliance. Those qualifications shall accompany the conclusion wherever it is communicated, including in summaries.

9.4 When a Quality Claim describes what the AI Service can do, it shall distinguish what the service has demonstrated from what is planned, intended by design, or promised in a contract. If the service demonstrated a capability only in a pilot, a test environment, or through a specific operating path, the Claim shall state that limit.

9.5 The issuing party shall maintain a record linking each material assertion and qualification in the Claim to the corresponding synthesized conclusion or Finding and, through that record, to the applicable requirement or criterion and supporting Evidence. The record shall identify the version of the Claim and the examination on which it is based.

9.6 Where a Quality Claim includes a Conformity Statement, the statement shall identify the requirements or criteria against which conformity was established and the scope and conditions to which the determination applies. The issuing party shall identify the party that performed the Conformity Assessment and state whether the determination was based on first-, second-, or third-party assessment.

9.7 Where a service change, new Evidence, or a revised Finding could materially affect continued reliance on a Claim, the issuing party shall review the Claim and document whether it remains applicable or requires qualification, revision, or withdrawal. Where the review changes the basis for reliance, the issuing party shall communicate the change to the affected recipients through the applicable communication channels and retain a record of the disposition.

Informative example. An organization deciding whether to use a benefits application service needs to know whether the service can confirm that the agency received an application. The Quality Claim states that agency receipt was verified for the examined cases using the original submission path. For the replacement path, the Claim states that queueing was verified and agency receipt remained unresolved. The qualification appears alongside the conclusion in the summary provided to the organization.

Bibliography

[1] Quality Institute. *Whole Quality of Complex Objects — How Objects Realize Their Intended Functions*. WQI_ARB1. First edition, November 2026.

[2] Digital Service Quality Initiative (DQI). *AI Service Conformity Assessment Guide*. Working draft.

[3] ISO/IEC 17000:2020, *Conformity assessment — Vocabulary and general principles*.

[4] ISO/IEC 22989:2022, *Information technology — Artificial intelligence — Artificial intelligence concepts and terminology*.

Copyright and Use Notice

© 2026 Digital Service Quality Initiative (DQI). All rights reserved.

This publication may be downloaded from the official DQI website and retained, printed, and used for personal, professional, educational, research, and internal organizational purposes. This Standard may be applied by reference in examinations, audit protocols, procurement requirements, and contractual documents, provided the document title, applicable edition or revision, and requirements being applied are identified. Such use does not authorize reproduction or modification of the publication or imply DQI endorsement, certification, or approval.

Brief quotations may be reproduced with appropriate attribution to the Digital Service Quality Initiative and identification of this publication and its applicable edition or revision.

Except as expressly permitted above, this publication may not be modified, adapted, translated, republished, redistributed, sold, or reproduced in another standard, specification, regulation, certification scheme, audit protocol, procurement requirement, contractual document, or similar instrument without prior written permission from DQI.

The official and current version of this publication is the version made available by the Digital Service Quality Initiative.

Document identification
DQI_AIS1 — AI Service Quality Standard — First Edition — November 2026.