

CIJARA GROUP LLC

AI Advisory | Risk | Governance
Charlotte, North Carolina | CijaraGroup.com

September 16, 2026

Razvan Amironesei and Jesse Dunietz

National Institute of Standards and Technology

By email: ai-standards+doczd@nist.gov

RE: Public comment on NIST AI 300-1 ipd, Guidance and Templates for Public-Facing AI Documentation (An AI Standards “Zero Draft”), initial public draft, July 2026.
<https://doi.org/10.6028/NIST.AI.300-1.ipd>

To the authors:

I am the Managing Principal of Cijara Group LLC, an AI advisory and risk governance firm in Charlotte, North Carolina. I write in my own capacity. This comment responds to your invitation for input on any aspect of the draft, and it uses the profile mechanism in Clause 6 to show one way the templates serve regulated deployers.

My background is second-line model risk and operational risk at a global systemically important bank (G-SIB), including Senior Model Risk Officer and Lead Operational Risk Officer roles. I disclose a commercial interest in Part 8.

One experience shapes this comment. In model validation, third-party vendor models were the hardest category to validate. Vendors withheld training data, equations, model logic, and code. Vendors are entitled to withhold. The cost was not the withholding. The cost was that the withholding was undeclared. The validation function could not scope compensating controls without negotiating the size of each gap with each vendor, one at a time. A documentation template fixes this at low cost. It asks a provider to declare, for any field left empty, whether the field is not applicable, not assessed, or withheld, and on what basis, without revealing the withheld content.

I have reviewed the publicly available comments on the September 2025 Extended Outline. That record includes documentation researchers and an industry trade association. I did not find a submission from the position of a second-line model risk function at a supervised institution. This comment is offered to fill that gap, and it refers to the prior round where the two intersect.

The submission has eight parts. Part 1 is a comment matrix in your template columns. Part 2 is a Regulated Financial Services Profile drafted in your Annex A format, which is the center of this submission. Part 3 is a one-page plain-language reader aid for the public audience Clause 4.2.2 names. Part 4 is a short machine-readable example toward Annex B. Part 5 answers each Note to Reviewers. Part 6 lists editorial corrections. Part 7 offers a pilot and a listening session. Part 8 contains disclosures.

A note on regulatory references. U.S. interagency model risk guidance was revised on April 17, 2026. Federal Reserve SR 26-2, OCC Bulletin 2026-13, and FDIC FIL-15-2026 supersede SR 11-7 and OCC Bulletin 2011-12, and OCC Bulletin 2026-13 also rescinded the Model Risk Management booklet of the Comptroller's Handbook. The revised guidance is principles-based, is expressly non-enforceable, and states that generative and agentic AI models are not within its scope. Every prudential reference in this submission is an informative crosslink explaining why a field is commonly needed by this audience. None is a mandate. This submission proposes no regulatory requirement, no enforcement, and no legislative vehicle. Consistent with your Foreword, every designation in the attached profile is a conformity criterion for voluntary adopters only.

I would welcome acknowledgment if you find the contribution useful, and I understand this submission and its attachments become public record.

Respectfully submitted,

Joshua M. Reh

Managing Principal, Cijara Group LLC

Charlotte, North Carolina

Joshua@CijaraGroup.com

<https://CijaraGroup.com>

Part 1. Comment matrix

Submitted by: Joshua M. Reh, Cijara Group LLC, Joshua@CijaraGroup.com. Page numbers are PDF pages of the published initial public draft counted from the cover. Line numbers are the draft's own running-text numbering. Tables in Clause 5 and Annex A are not line-numbered, so those rows cite the table and field ID.

No.	Type	Page and line	Clause or field	Comment and rationale	Suggested change
1	Technical	p21, line 591	Clause 5.1, general requirements	The three designations govern the provider's obligation to populate a field. None tells a reader what an absent field means. An absent field may be not applicable, not assessed, or assessed and withheld. Those carry different risk treatments for a deployer. In second-line validation, resolving that ambiguity required bilateral negotiation with each vendor. Declaring an omission discloses nothing withheld and is compatible with Clause 4.2.4 judiciousness, Clause 3.4.3 proprietary information protection, and Clause 3.3.2 malicious use.	For NIST's consideration: permit, for any field left unpopulated, an optional disclosure-status marker with one value from a closed list, NOT APPLICABLE, NOT ASSESSED, or WITHHELD, and where WITHHELD, a basis from a closed list, proprietary, security, legal, or privacy. The marker records only the fact and basis of omission, never the content. If NIST prefers not to amend Clause 5, the same construct is demonstrated as a neutral element of the profile in Part 2.
2	Technical	p47 and p48, table rows	Annex A, Table A.3.1, Fields 5.2-(*) and 6.1.1-(*)	Both fields state that where a referenced dataset cannot be publicly disclosed, the artifact should still describe characteristics readers likely need to understand. That language does not appear in the September 2025 Extended Outline and appears to respond to prior-round feedback. Both then add "to the extent consistent with the provider's objectives." That qualifier returns the decision to the party whose objective may be to disclose nothing. An artifact describing no characteristics remains conformant. The recommendation cannot fail and therefore cannot bind.	Replace the qualifier with: "Where the referenced dataset cannot be publicly disclosed, the documentation artifact should describe the characteristics necessary for a downstream implementer to assess suitability and scope compensating controls, and should indicate which characteristics have been withheld."
3	Technical	p48 and p50, table rows	Annex A, Table A.3.1, Fields 6.1 and 6.5-(*)	Field 6.1 requires only which methods are automated versus human. There is no field for the configuration under which a reported score was produced. Identical weights produce materially different results under different harnesses. In one frontier release in September 2026, the same model and weights scored 62.7 percent on a public reasoning benchmark under a provider-neutral harness and 99.9 percent under a	Add a Recommended subfield under Field 6.5-(*): Evaluation Configuration. The harness or scaffold used, tool access granted, number of attempts permitted, and any effort, compute, or reasoning-budget setting applied to the reported result.

No.	Type	Page and line	Clause or field	Comment and rationale	Suggested change
				provider-supplied harness. Without configuration, a reported result is neither reproducible nor comparable, which defeats Clause 3.3.1.2 suitability assessment and the Clause 4.2.1 correctness quality.	
4	Technical	p50, table row	Annex A, Table A.3.1, Field 6.5-(*). Prior-round item not adopted.	The prior round recommended that NIST encourage disclosure of whether training and test data overlap, or whether that overlap was evaluated. The initial public draft contains no such field. The consequence is measurable. In the same frontier release, a headline capability score of 100 percent fell to 39.0 percent when the same capability was measured on material postdating the training data cutoff. Both figures were published by the provider, in different documents, with no field placing them side by side. The prior round's related request on robustness versus reliability was adopted as Field 6.5-(*).3, so this asks for the same treatment of a second attribute of a reported result.	Add a Recommended subfield under Field 6.5-(*): Contamination Control. Whether the evaluation data postdates the training data cutoff or whether the overlap was otherwise assessed, the decontamination method, and where a contamination-controlled variant exists, its result reported alongside the headline result.
5	Technical	p49, table row	Annex A, Table A.3.1, Field 6.2 Third-Party Evaluation	The field is Optional and captures identity, affiliation, nature, scope, and access conditions. It omits the four attributes that determine how much weight a reader can place on an external evaluation. An evaluation commissioned by the provider, run over a short window, funded by the provider, and publishable only through the provider's own release is a different evidence source from an arms-length evaluation with an independent right to publish. In one recent release three external evaluators occupied three different postures on these attributes, and the artifact gave a reader no standard field to tell them apart.	Elevate Field 6.2 to Conditionally Recommended, recommended where any third-party evaluation is referenced anywhere in the artifact. Add four subfields: Commissioning Relationship (commissioned or arms-length), Funding Relationship, Evaluation Duration or Window, and Right to Publish Independently.
6	Technical	p50, table row	Annex A, Table A.3.1, Field 6.6 Risk and Impact Analyses	The designation trigger is retrospective. The field is recommended only where negative impacts have previously been publicly documented and traced to this model or similar models. A model	Add a forward-looking condition. Recommended where negative impacts have previously been documented, or where the provider's own evaluations identify a capability the provider

No.	Type	Page and line	Clause or field	Comment and rationale	Suggested change
				introducing a genuinely novel capability has no such history, so the field is not recommended in the case where risk analysis matters most.	assesses as materially exceeding that of prior comparable models.
7	Technical	p51, table row	Annex A, Table A.3.1, Field 7.3 Change Log	The field is Optional and describes changes after deployment without distinguishing improvements from declines. A provider has a standing incentive to narrate the former and omit the latter. A deployer treats a measured decline against a prior version as a control trigger. In one recent release, a decline in a documented monitoring property was disclosed while a benchmark result trailing three named competitors was present in the provider's own table and absent from its announcement prose.	Elevate Field 7.3 to Recommended and add to the guidance column: "Where a documented measure declines relative to the prior version, the change log should state the decline."
8	Technical	p53, table row	Annex A, Table A.3.1, Field 8.2.1 Documentation Quality Control	The guidance already contemplates a record of when and by whom accuracy was reviewed but leaves it as an example. Named accountability for accuracy is the lowest-cost quality control available. This is distinct from the Clause 6 attestation of conformity to a template or profile and should be kept separate in the text.	Make the reviewer record an explicit Recommended subfield: Reviewer of Record. The name or role of the individual who reviewed the artifact for accuracy, and the date of review.
9	General	p9 and p10, lines 129 to 153	Clause 2, definition of documentation artifact, and the associated Note for Reviewers	Responding to the note on terminology. "Artifact" reads as jargon to the public audience Clause 4.2.2 names, and in governance and audit practice it often connotes an incidental byproduct rather than a controlled deliverable.	Adopt "documentation record." It is plain, it aligns with the established use of "record" in ISO management-system standards, and it signals a maintained and versioned deliverable. If "artifact" is retained, add a one-line gloss on first use stating it is a controlled and versioned deliverable.
10	General	p10, line 164 and the following note	Clause 2, definition of AI model, and the Note for Reviewers on scope	The note asks how model documentation scope could expand without encompassing entire systems, and states NIST has not identified a principled basis for scoping some components in and others out. The prior round challenged the exclusion of system-level documentation but proposed no line. A measurement-conditions test supplies one. An evaluation result	Scope in the conditions under which the model was measured, including whether a reported result was obtained with safeguards active or inactive. Scope out deployed safeguard components. State the test in the note to Clause 2 so future profiles and revisions apply it consistently. The test is self-limiting, admits measurement

No.	Type	Page and line	Clause or field	Comment and rationale	Suggested change
				obtained with output filters or guardrails disabled is a property of the model object, because it describes what the architecture and parameters do. A deployed guardrail stack is a property of the system, because a different deployer configures it differently or not at all.	conditions and no components, and is compatible with treating full system documentation as future work.
11	General	p22, line 618 note, and p30, line 697 note	Clause 5.1 Note to Reviewers on examples, and Annex A Note to Reviewers on pilots	Responding to the notes on examples and pilots. A deployer's most common use of a standard like this is to hand it to a vendor as a procurement specification. A complete worked example serves that purpose better than inline illustrations, which would enlarge the templates and reduce the comprehensibility Clause 4.2.2 seeks.	Provide fully populated example artifacts in a separate annex rather than inline. The machine-readable example in Part 4 is offered as a starting point toward Annex B. The pilot offer is in Part 7.
12	General	p28 and p29, lines 628 to 696	Clause 6, Profiles	The profile mechanism is well designed and is the correct home for sector-specific and risk-tier-specific requirements. No change is sought. It is noted because it is the vehicle for Part 2 and explains why Comment 1 is offered for NIST's consideration at Clause 5 rather than as a demand: a profile can require a field, but nothing in the current draft lets a profile give meaning to a field's absence.	Retain Clause 6 as drafted.
13	Editorial	p5, lines 28 to 30	Contents, Clauses 4.2.1 to 4.2.3	Contents lists 4.2.1 comprehensibility, 4.2.2 informativeness, 4.2.3 correctness. The body at lines 385, 392, and 400 orders them correctness, comprehensibility, informativeness.	Align the Contents to the body ordering, or renumber the body.
14	Editorial	p15, line 352	Clause 3.4.5 Trust	The text refers to artifact qualities in Clause 4.3. Artifact qualities are in Clause 4.2.	Change the cross-reference to Clause 4.2.
15	Editorial	p17, line 444	Clause 4.2.6 Artifact interoperability	The text refers to continuous documentation processes in Clause 4.2.6, the clause making the reference. The intended target appears to be Clause 4.3.6.	Change the cross-reference to Clause 4.3.6.
16	Editorial	p32, line 723 and p43, line 730	Annex A.2 and A.3 introductory sentences	Both state that fields are drawn from Clause 5.3. The dataset profile in A.2 draws from Clause 5.2.	Correct the A.2 reference to Clause 5.2.
17	Editorial	p43 and p44, table rows	Annex A, Table A.3.1, Fields 1.6	Field 1.6 refers to 1.7.1 and 1.7.3, which do not exist, and appears to	Renumber the identifying-descriptor subfields into a

No.	Type	Page and line	Clause or field	Comment and rationale	Suggested change
			and 1.7, and field numbering	mean 1.6.1 and 1.6.3. Field 1.7 refers to 1.8.1, which appears to be 1.7.1. There is no Field 1.4, so the model profile is misaligned with the dataset profile, where 1.4 is Point of Contact.	continuous sequence aligned with the dataset profile, and repair the internal references.
18	Editorial	p44, table row	Annex A, Table A.3.1, Field 1.12 Documentation Version Identifier	The contents of the Description and Designation columns are transposed.	Transpose the two cells.
19	Editorial	p25 to p27 and p44 to p52, table rows	Clause 5.3 Table 2 and Annex A Table A.3.1	"Dataset" appears where "model" is intended in Table 2 Fields 2, 3, and 8, and in model profile Fields 2.1, 3.3, 7.4, 7.9, and 8. The same substitution was flagged in the prior round on the Extended Outline and persists in this draft in at least eight places.	Replace "dataset" with "model" at each occurrence in the model template and model profile.
20	General	p3 to p4, lines 80 to 104	Foreword, use of shall and requirement	Confirming the framing this submission adopts. The Foreword states that shall and requirement denote conformity only, not regulation. Every designation proposed in Part 2 is offered on that basis.	No change requested.

Part 2. Regulated Financial Services Profile (candidate, Clause 6)

Scope

This profile specializes the Clause 5 model template, and two fields of the dataset template, for regulated financial institutions and the third-party model and data providers that supply them. It uses only the Clause 6 profile mechanism. It elevates designations and adds subfields. It changes nothing in Clause 5 and contradicts nothing in it. Non-bank regulated deployers may adopt it. This profile is deliberately narrow. It adds no fields for information security, adversarial testing, environmental impact, fairness outside credit and other consumer-financial decisions, or system-level deployment controls. Those subjects fall outside the measurement-conditions boundary described in Part 5, or are better served by a separate profile. The omissions are scope decisions, not oversights.

Nature of the regulatory references

Every regulatory reference below is an informative crosslink explaining why a documentation field is commonly needed by this audience. None is a mandate. This profile creates no legal obligation. Adoption is voluntary. Consistent with the draft's Foreword, shall, required, and recommended denote conformity criteria for voluntary adopters only.

Two limits, stated plainly

First, U.S. interagency model risk guidance was revised on April 17, 2026. Federal Reserve SR 26-2 supersedes and replaces SR 11-7. OCC Bulletin 2026-13 and FDIC FIL-15-2026 are the parallel issuances. The revised guidance is principles-based, is expressly non-enforceable, is expected to be most relevant to banking organizations with over \$30 billion in total assets, and states that generative AI and agentic AI models are not within its scope. Where this profile serves generative or agentic models, the prudential crosslinks do not apply to those model classes and are offered by analogy only. The evaluation-configuration and contamination-control fields rest on measurement science, not on supervisory authority.

Second, OCC Bulletin 2026-13 rescinded the Model Risk Management booklet of the Comptroller's Handbook. Field-level detail below is therefore anchored to surviving principles, to the 2023 Interagency Guidance on Third-Party Relationships (OCC Bulletin 2023-17, Federal Reserve SR 23-4, FDIC FIL-29-2023), to the Basel Committee's Principles for Effective Risk Data Aggregation and Risk Reporting (BCBS 239), and to CFPB Circulars 2022-03 and 2023-03.

Disclosure-status convention

For any field this profile designates Required or Recommended that a provider does not populate, a conforming artifact should record a disclosure status of NOT APPLICABLE, NOT ASSESSED, or WITHHELD. Where WITHHELD, the artifact should record a basis of proprietary, security, legal, or privacy.

This convention discloses only the fact and basis of omission. It never requires disclosure of withheld content. It mirrors the informativeness-versus-judiciousness trade-off in Clause 4.4.1.

Model template additions and elevations

IDs follow Annex A Table A.3.1. New subfields take the next available number under the parent.

ID	Field name	Description	Designation	Additional guidance and requirements
1.13	Inventory Identifier	Stable identifier linking the model version to the adopter's model inventory record.	Recommended	Enables the adopter to tie the artifact to inventory, ownership, and validation status. Informative: SR 26-2 and OCC 2026-13 retain a firmwide inventory as

ID	Field name	Description	Designation	Additional guidance and requirements
				sound practice.
2	Intended Use	Applications for which the model was designed, and known unsuitable uses.	Recommended (elevated from Optional)	Record intended tasks, application domains, and uses the model is known to be unsuitable for. Correct the base-template reference from dataset to model.
4	Design	Architecture, technical assumptions, and key assumptions.	Recommended (elevated from Optional)	See 4.4 and 4.5.
4.4	Conceptual Soundness Basis	Theory, evidence, or published practice supporting the design's fitness for the stated use, and alternatives considered.	Recommended	Informative: SR 26-2 and OCC 2026-13 treat conceptual soundness as a component of sound development and validation.
4.5	Key Assumptions	Assumptions whose failure would most degrade performance, and the conditions under which they hold.	Recommended	Written so that a party unfamiliar with the model understands its limits, consistent with the documentation standard retained in SR 26-2.
5	Training	Protocols and data used to optimize the model.	Recommended (elevated from Optional). Conditionally Required where the provider represents the model as fit for a regulated financial use.	See 5.3.
5.3	Training Data Characteristics	Characteristics a downstream implementer needs to assess suitability and scope compensating controls, and an indication of which characteristics are withheld.	Recommended in all cases. Conditionally Required where any training dataset cannot be publicly disclosed.	Apply the disclosure-status convention. Structure per the dataset template in Clause 5.2. Informative: BCBS 239 Principles 3 and 4. State whether any portion of the training data was generated by another model, and the approximate proportion, so a downstream implementer assesses data quality and the risk of degradation from training on model output.
6	Evaluation	Evidence of quality and fitness for use.	Recommended (elevated from Optional)	See 6.2, 6.5, 6.6, and 6.7 through 6.9.
6.2	Third-Party Evaluation	Evaluation conducted by parties external to the developer, and the conditions under which it was performed.	Conditionally Recommended, recommended where any third-party evaluation is referenced	Add subfields 6.2.1 Commissioning Relationship (commissioned or arms-length), 6.2.2 Funding Relationship, 6.2.3 Evaluation Duration or Window, 6.2.4 Right to Publish Independently.
6.5	Robustness and	Behavior under adverse, shifted,	Recommended	Robustness and reliability are different

ID	Field name	Description	Designation	Additional guidance and requirements
	Reliability	or out-of-distribution conditions, and consistency of performance over repeated use, reported as distinct properties.	(elevated from Optional)	properties and should not be collapsed into one statement. Consistent with subfield 6.5-(*) ³ adopted after the prior round.
6.6	Risk and Impact Analyses	Analyses of known vulnerabilities, risks, and impacts, with supporting evidence.	Recommended. Conditionally Required where the model introduces a novel or materially expanded capability.	Broadens the base trigger from retrospective only to include a forward-looking condition. Populating this field can itself be a risk treatment, consistent with Annex A.1.
6.7	Evaluation Configuration	Conditions under which each reported result was produced: harness or scaffold, tool access, number of attempts, effort or compute setting.	Recommended	A score without its configuration is neither reproducible nor comparable. Evidence in Part 1, Comment 3.
6.8	Contamination Control	Whether evaluation items postdate the training cutoff or overlap was otherwise assessed, the method used, and any contamination-controlled result reported beside the headline.	Recommended	Evidence in Part 1, Comment 4. Responds to the prior-round request NIST did not adopt.
6.9	Fair-Lending and Disparate-Impact Testing	Testing methodology and results for protected-class outcomes, and whether the model supports specific and accurate reasons for adverse action.	Conditionally Recommended, recommended where the model informs credit or other consumer-financial decisions	Informative: ECOA and Regulation B. CFPB Circular 2022-03 states that ECOA and Regulation B do not permit creditors to use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions. CFPB Circular 2023-03.
7	Maintenance and Monitoring	Processes and responsibilities after release.	Recommended (elevated from Optional)	See 7.3 and 7.10.
7.3	Change Log	Changes made after deployment, including declines relative to the prior version.	Recommended (elevated from Optional)	Where a documented measure declines relative to the prior version, the change log shall state the decline. Where a change materially alters intended use, measured performance, or the risk profile, the documentation record shall be reissued rather than amended in place, and the superseded version shall remain retrievable.
7.10	Performance Thresholds and Triggers	Quantitative thresholds on monitored measures that prompt review, and the escalation path when breached.	Recommended	Informative: SR 26-2 and OCC 2026-13 ongoing monitoring.
8	Governance	Organizational practices applied	Recommended	See 8.6 through 8.8. Correct the base-

ID	Field name	Description	Designation	Additional guidance and requirements
		over the model lifecycle.	(inherited)	template reference from dataset lifecycle to model lifecycle.
8.6	Roles	Named roles for developer, owner, and independent validator, with reporting-line independence of the validator stated.	Recommended	Informative: SR 26-2 and OCC 2026-13 effective challenge and independence.
8.7	Reviewer of Record	Name or role of the individual who reviewed the artifact for accuracy, and the date of review.	Recommended	Distinct from a Clause 6 attestation of conformity, which is a claim about the artifact's relationship to a template or profile, not about who checked its accuracy.
8.8	Third-Party Access and Contractual Provisions	What the provider makes available to a downstream validator (documentation, evaluation access, escrow, audit rights) and what is withheld, with basis.	Conditionally Recommended, recommended where the model is supplied by a third party	Apply the disclosure-status convention. Informative: 2023 Interagency Guidance on Third-Party Relationships (OCC 2023-17, SR 23-4, FIL-29-2023). SR 26-2 and OCC 2026-13, third-party section.

Dataset template additions and elevations

IDs follow Annex A Table A.2.1.

ID	Field name	Description	Designation	Additional guidance and requirements
4.7	Data Flows	Chain of custody from source system to model input, including transformations and reconciliation points.	Recommended (elevated from Optional)	Source-to-input lineage sufficient to trace an input to its origin. Informative: BCBS 239 Principles 2 and 3.
7.6	Field-Level Definitions	A data dictionary defining each field a downstream implementer must interpret.	Recommended	Apply the disclosure-status convention where definitions are withheld. Informative: BCBS 239 Principle 2.

Proportionality

Consistent with the risk-based approach in SR 26-2, an adopter may apply this profile proportionately, populating more fields for higher-materiality models and fewer for lower-materiality ones. A community institution with no dedicated model risk staff may reasonably treat Fields 1.13, 6.7, 6.8, and 8.8 as the minimum useful set to request from a core vendor.

Part 3. Plain-language reader aid

For the general public audience named in Clause 4.2.2. Intended to accompany any artifact conforming to the profile in Part 2.

What this document is. A model record tells you what an AI model is, what it is for, how well it works, and who is responsible for it. This profile adds a few plain questions for models used in banking and finance.

What to look for:

- **What is it, and which version.** Can you tell exactly which model you are looking at.
- **What is it for.** And what is it not for.
- **How well does it work, and under what test conditions.** A score means little without the setup that produced it.
- **What changed in the latest version.** Including anything that got worse, not only what got better.
- **Who checked it.** A name or a role, and a date.
- **If a section is blank, why.** Not applicable, not tested, or withheld. If withheld, is it for a proprietary, security, legal, or privacy reason.

The last point is the heart of this comment. A blank you understand is far more useful than a blank you cannot.

Part 4. Machine-readable example toward Annex B

This example reuses existing conventions so it is easy to adopt: model-card YAML front matter as used on public model hubs for models, and the MLCommons Croissant JSON-LD format, built on schema.org, for datasets. The disclosure-status convention is expressed as a simple optional per-field object. Values are illustrative.

```
# Model documentation record (illustrative)
# Aligned to the Clause 5 model template and the Part 2 profile
identifying_descriptors:
  model_name: "Example Credit Risk Model"
  version: "2.3.0"
  inventory_identifier: "MRM-04821"
intended_use:
  tasks: ["probability-of-default estimation"]
  unsuitable_uses: ["real-time fraud blocking"]
design:
  conceptual_soundness_basis: "logistic regression on validated risk factors"
  key_assumptions: ["macro conditions within observed range"]
training:
  training_data_characteristics:
    status: "WITHHELD"
    basis: "proprietary"
evaluation:
  evaluation_configuration:
    harness: "provider-neutral v1.2"
    tool_access: "none"
    attempts: 1
    effort_setting: "standard"
  contamination_control:
    items_postdate_training_cutoff: true
    method: "temporal holdout"
  third_party_evaluation:
```

```
    status: "NOT APPLICABLE"
  fair_lending_testing:
    status: "NOT ASSESSED"
  maintenance_and_monitoring:
    change_log:
      improvements: ["recalibrated PD curve"]
      declines: ["AUC down 0.4 points on segment C"]
      performance_thresholds: ["review if PSI exceeds 0.25"]
  governance:
    roles:
      developer: "Model Development"
      owner: "Retail Credit"
      independent_validator: "Model Risk Management"
    reviewer_of_record:
      role: "Senior Model Risk Officer"
      date: "2026-08-30"
    third_party_access:
      status: "NOT APPLICABLE"
```

For a dataset record, the same fields map onto Croissant JSON-LD, with `data_flows` and `field_level_definitions` carried as additional properties and the same status and basis objects for withheld content.

Part 5. Responses to the Notes to Reviewers

Terminology for documentation artifact (Clause 2)

Adopt “documentation record.” It is plainer for public readers, it aligns with the established use of “record” in ISO management-system standards, and it signals a maintained and versioned deliverable. If “artifact” is retained, add a one-line gloss on first use.

Expanding model documentation scope without encompassing systems (Clause 2)

Use a measurement-conditions test. Scope in the conditions under which the model was measured, including results obtained with safeguards disabled, because those describe the model object. Scope out deployed safeguard components such as guardrail classifiers, because those describe the system. This supplies the principled basis the note says NIST has not found. It draws the boundary at what the documentation describes, not at where a component sits. It is self-limiting and compatible with treating full system documentation as future work, which the prior round requested.

Inline examples versus an annex (Clause 5)

Use a separate annex of fully populated examples. Inline examples would enlarge the templates and reduce comprehensibility. Part 4 is offered as a starting point.

Changes to default profile designations (Annex A)

Part 2 elevates Design, Training, Evaluation, and Maintenance and Monitoring, and explains per field how the change improves outcomes on the Clause 3 dimensions of suitability assessment, accountability, and trust, and supports the Clause 4 qualities of informativeness and correctness. The elevations are confined to a defined adopter context and do not burden general use.

Clarifying field content so documenters populate fields reliably (Annex A)

Two practices help. First, the disclosure-status convention removes ambiguity about blanks. Second, for high-value fields, prefer the prescriptive framing the Annex A note floats, stating what a field shall, should, or may contain, over a pure noun-phrase definition. The guidance column in Part 2 illustrates this.

Whether fields are sufficiently distinguished (Annex A)

Mostly yes. The recurring confusion is between Usage Rights and Restrictions and Governance. The dataset template addresses it by cross-reference in both directions. The model template Field 3 cross-references Field 8 but Field 8 does not cross-reference Field 3. Add the reverse reference.

Invitation to pilot (Clause 6)

Accepted. See Part 7.

Part 6. Editorial corrections

Offered as a courtesy. The final item was raised in the prior round and persists.

1. Contents (lines 28 to 30) orders 4.2.1 comprehensibility, 4.2.2 informativeness, 4.2.3 correctness. The body (lines 385, 392, 400) orders correctness, comprehensibility, informativeness.
2. Clause 3.4.5 (line 352) cites Clause 4.3 for artifact qualities. Should be Clause 4.2.
3. Clause 4.2.6 (line 444) cites itself for continuous documentation. Should be Clause 4.3.6.
4. Annex A.2 (line 723) and A.3 (line 730) both cite Clause 5.3. A.2 should cite Clause 5.2.
5. Model profile Field 1.6 refers to nonexistent 1.7.1 and 1.7.3. Field 1.7 refers to 1.8.1. There is no Field 1.4. Renumber and repair references.

6. Model profile Field 1.12 has Description and Designation columns transposed.
7. “Dataset” appears where “model” is intended in Table 2 Fields 2, 3, and 8, and in model profile Fields 2.1, 3.3, 7.4, 7.9, and 8. Flagged in the prior round. Persists in at least eight places.

Part 7. Offer to pilot and to host a listening session

I offer to pilot the Regulated Financial Services Profile and the base templates against published model documentation from more than one provider and to share what I learn, consistent with the Note to Reviewers inviting pilots. I have already applied a prototype profile to one recently released frontier model using only that provider's public documentation, populating a disclosure status for every field. The exercise surfaced an evaluation domain with no published result and no statement of whether it was assessed, a declared monitoring limitation that was not defined, and an external evaluation whose duration and publication rights were not stated. I am willing to share that pilot with NIST in whatever form is most useful.

I also offer to host a listening session convening model risk, vendor risk, and data governance practitioners from regulated institutions to give NIST verbal feedback, as the program page invites.

Part 8. Disclosures

Commercial interest

I am the Managing Principal of Cijara Group LLC, an AI advisory and risk governance firm. I advise institutions on model risk and operational risk. I hold no federal contract or subcontract related to this rulemaking, and no financial interest in its outcome. If a documentation standard of this kind is adopted, my firm could benefit from advising on its implementation. I disclose this so the comment is read with that interest in view. My recommendations are offered as neutral, unbranded field designations, not as a product.

AI assistance

As requested in the Note to Reviewers, I confirm that an AI assistant was used in preparing this submission. It was used to review the initial public draft against a documentation framework I had already developed, to review the prior-round comment record, to identify internal reference errors, to check citations, and to draft and structure this document. The professional experience in the cover letter, the assessment of the disclosure-status gap, the scope test in Part 5, and all recommendations are my own and reflect direct work on the deployer side of vendor model validation. I verified the regulatory and technical claims against primary sources and take responsibility for the content.