

DP23 – Universal Participation & Linguistic Interoperability

The Conditions for Shared Global Sensemaking

1. Purpose of This Draft

This ML-Draft articulates **Desirable Property 23 (DP23)** as a foundational condition for universal participation in the Meta-Layer.

It expands multilingual support beyond interface localization and machine translation to encompass:

- linguistic inclusion
- cross-language semantic interoperability
- cultural and contextual integrity
- community stewardship of language systems
- equitable participation in global issues
- traceable multilingual collective sensemaking
- visible feedback between public contribution and institutional response

DP23 responds to several overlapping needs:

- The need for people to participate without first adopting a dominant language
- The need for participants to understand contributions made across linguistic boundaries
- The need to preserve meaning, ambiguity, tone, cultural context, and local knowledge during translation
- The need to prevent dominant-language communities from disproportionately shaping shared understanding
- The need for AI systems to reason across languages without treating translation into English or another dominant language as the only path to comprehension
- The need for communities to inspect, challenge, correct, and govern multilingual AI outputs
- The need for participants to see how their contributions influenced collective understanding, recommendations, decisions, and action

This draft is intended to guide implementation, governance design, evaluation, community stewardship, and future ML-RFC development.

2. Problem Statement: Why Translation Alone Is Not Enough

For most of the Web's history, language inclusion has been treated as a content-delivery problem. A page is written in one language. A translation system converts it into another. If the translated text is readable, the problem is assumed to have been solved.

This assumption no longer holds.

Translation can increase access to words while still distorting meaning, flattening cultural context, suppressing minority interpretations, and reinforcing the dominance of languages with the strongest technical and commercial support.

A person may technically be able to read translated content while remaining unable to participate equitably in the process that interprets, organizes, summarizes, and acts upon it.

DP23 begins from a different premise:

Linguistic access is not achieved when words are merely translated. It is achieved when people can participate, be understood, inspect interpretation, challenge distortion, and influence shared outcomes without surrendering their language or cultural context.

Language is therefore not an optional accessibility feature of the Meta-Layer.

It is part of the participation boundary.

If language determines whose ideas are legible, searchable, classifiable, summarized, or acted upon, language also determines whose knowledge counts.

2.1 The Limits of Translation-Centric Inclusion

Translation-centric systems answer a narrow question:

How can this text be rendered in another language?

They do not adequately address:

- whether the translation preserves the intended meaning
- whether uncertainty and ambiguity remain visible
- whether culturally specific concepts have equivalents
- whether dominant terminology has displaced local framing
- whether translated contributions remain connected to their originals
- whether people can challenge machine interpretations
- whether low-resource languages receive comparable system quality

- whether multilingual contributions affect summaries and decisions equitably
- whether readers know when meaning has been inferred rather than translated
- whether contributors can see what happened after they participated

A fluent translation can still be epistemically misleading.

A grammatically correct summary can still erase disagreement.

A multilingual interface can still reproduce linguistic hierarchy if the underlying reasoning, clustering, moderation, and governance systems privilege a dominant language.

2.2 Structural Failure Modes in Today's Web

Several structural conditions prevent meaningful global participation:

- **Language silos:** Conversations about the same issue remain separated across linguistic communities.
- **Dominant-language mediation:** Contributions are translated into a dominant language before they become searchable, classifiable, or institutionally visible.
- **Semantic flattening:** Terms with cultural, legal, spiritual, ecological, historical, or political significance are reduced to approximate equivalents.
- **Unequal model quality:** High-resource languages receive better translation, moderation, retrieval, and summarization than underrepresented languages.
- **Invisible uncertainty:** Machine-generated translations appear authoritative even when multiple interpretations are plausible.
- **Lost provenance:** Readers see translated or summarized claims without an inspectable connection to the original contribution.
- **Extractive participation:** Institutions collect public input without showing contributors how it was grouped, interpreted, used, or acted upon.
- **Majority-language amplification:** Views expressed in widely supported languages are easier to search, quote, summarize, and elevate.
- **Cultural decontextualization:** Local knowledge is separated from the practices, histories, relationships, and worldviews that give it meaning.
- **Classification bias:** Categories developed in one linguistic or cultural setting are applied globally as if they were neutral.
- **Synthetic consensus:** AI systems combine multilingual disagreement into a smooth summary that suggests greater agreement than the corpus supports.
- **Missing constituency blindness:** Systems summarize the people who participated without identifying who was absent, excluded, or unable to participate.

These are not isolated translation errors.

They are systemic properties of digital environments that treat language as a formatting problem rather than a condition of civic agency.

2.3 DP23 as a Shift in Framing

DP23 reframes multilingual participation along several axes:

- From **translation** to **linguistic interoperability**
- From **access to text** to **access to participation**
- From **dominant-language normalization** to **language-plural reasoning**
- From **machine authority** to **community-governed interpretation**
- From **content localization** to **cultural and contextual integrity**
- From **consultation** to **traceable influence**
- From **international audiences** to **shared global issue participation**
- From **one-way communication** to **multilingual collective sensemaking**

Rather than asking only whether content can be translated, DP23 asks:

- Under what conditions can participation remain equitable across languages?
 - How is meaning preserved across linguistic and cultural contexts?
 - Who can challenge interpretations?
 - How are multilingual contributions compared and synthesized?
 - How are missing perspectives made visible?
 - How can contributors trace their influence on collective outcomes?
-

3. Threats and Failure Modes

The threat classes below are non-exhaustive. They describe the risk context DP23 exists to address rather than a complete adversarial catalogue.

DP23 is not defined in opposition to any particular language, institution, country, or technology.

It responds to recurring failure modes that emerge whenever linguistic asymmetry intersects with governance, AI mediation, knowledge production, and large-scale participation.

3.1 Linguistic Exclusion

People may be unable to participate because:

- the interface is unavailable in their language
- contribution rules require a dominant language

- moderation does not support their language
- translation quality is too poor for meaningful use
- required terminology has no agreed equivalent
- their writing system is not rendered correctly
- voice, literacy, disability, or device constraints prevent text-based participation
- institutional staff cannot review their contributions

DP23 treats these conditions as participation failures, not merely incomplete localization.

3.2 Meaning Distortion

Translations may change:

- certainty into speculation
- speculation into certainty
- criticism into aggression
- formality into informality
- sacred language into ordinary terminology
- collective responsibility into individual responsibility
- relational concepts into transactional concepts
- culturally specific identities into externally imposed categories

Distortion may occur without malicious intent.

Machine systems often select a plausible equivalent where the source carries several possible meanings.

DP23 therefore requires ambiguity, uncertainty, and contested terminology to remain visible rather than being silently resolved.

3.3 Linguistic and Cultural Capture

A multilingual system can appear inclusive while remaining controlled by:

- dominant-language taxonomies
- centralized translation providers
- institutional glossaries
- commercially supported model languages
- professional experts disconnected from affected communities
- political actors able to define preferred terminology
- platform owners who determine which interpretations are visible

DP23 counters capture by requiring plural terminology, community participation, contestability, version history, and visible governance authority.

3.4 Epistemic Extraction

Participants may be invited to contribute experiences, knowledge, and cultural insight without receiving:

- attribution
- acknowledgment
- compensation where appropriate
- visibility into how their contribution was used
- access to resulting analysis
- a response from responsible institutions
- an opportunity to correct interpretation
- meaningful influence over outcomes

In such cases, multilingual participation becomes a mechanism for extracting knowledge rather than distributing civic power.

DP23 requires visible feedback loops between contribution, synthesis, recommendation, responsibility, response, and action.

3.5 Automated Misclassification

AI systems may incorrectly classify a contribution because:

- sarcasm is culturally specific
- political terminology is context-dependent
- a local idiom is translated literally
- categories do not map cleanly between cultures
- code-switching is mistaken for incoherence
- dialect is treated as error
- Indigenous or oral knowledge is forced into institutional categories
- disagreement with a summary is classified as hostility

Misclassification can affect moderation, visibility, reputation, clustering, and inclusion in summaries.

Participants must therefore be able to correct classifications and inspect their consequences.

3.6 Synthetic Global Consensus

Large multilingual corpora create pressure to produce concise summaries.

This can cause systems to:

- collapse differences into shared themes
- treat repetition as representativeness
- omit low-frequency perspectives
- hide unresolved contradictions
- merge culturally distinct proposals
- overstate agreement
- mistake absence of participation for absence of concern

DP23 does not prohibit synthesis.

It requires synthesis to remain traceable, plural, contestable, and explicit about missing participation.

3.7 Cross-Language Manipulation

Bad actors may exploit multilingual systems through:

- coordinated messaging across languages
- translation poisoning
- strategic mistranslation
- false claims of cultural authority
- glossary capture
- fabricated constituency representation
- AI-generated submissions at synthetic scale
- inconsistent narratives targeted at separate language communities
- deliberate exploitation of weak moderation languages

DP23 works alongside DP1, DP11, DP12, DP13, DP14, DP15, and DP18 to make such behavior attributable, visible, challengeable, and governable.

4. Core Principle

Language is not merely descriptive at the interface layer.

It shapes:

- access

- discoverability
- comprehension
- classification
- credibility
- participation
- representation
- collective memory
- institutional responsiveness

Language is therefore a core dimension of civic infrastructure.

If meaning cannot move across linguistic contexts without losing provenance, uncertainty, and contestability, global participation collapses into selective interpretation.

4.1 Core Principle of DP23

DP23 establishes universal participation and linguistic interoperability as foundational conditions for shared collective sensemaking in the Meta-Layer.

Every person should be able to contribute in a language they can meaningfully use, understand the relevant corpus in a language they can meaningfully understand, inspect how interpretation occurred, challenge errors or omissions, and trace how their participation affected collective outcomes.

This principle has several direct implications.

Language choice is an expression of agency

Participants should not be required to adopt a dominant language to be considered legible, credible, or relevant.

Original expression remains canonical

Translations, classifications, embeddings, clusters, summaries, and recommendations are derived representations.

They must never silently replace the source contribution.

Translation is contextual, not merely lexical

Meaning depends on issue, community, culture, history, speaker intent, and situational context.

Translation systems must be able to use issue-specific and community-governed terminology.

Multilingual AI remains uncertain and contestable

AI must not conceal ambiguity behind fluent output.

Participants and language stewards must be able to challenge translations, classifications, clusters, and summaries.

Shared understanding does not require linguistic uniformity

Participants may inhabit the same issue space while contributing and reading in different languages.

The Meta-Layer should support convergence of attention without forcing convergence of language.

Participation includes a right to trace influence

People should be able to see whether their contribution was included, grouped, cited, challenged, incorporated into a recommendation, or left without documented influence.

Global participation remains issue-specific

DP23 does not assume every issue requires participation from every person.

It enables appropriately scoped global participation where an issue crosses jurisdictions, populations, ecosystems, generations, or shared human interests.

Absence must remain visible

A multilingual corpus cannot be treated as globally representative merely because it includes several languages.

Missing linguistic, cultural, geographic, professional, generational, and affected constituencies must be surfaced.

5. Primary Mechanisms and Structural Conditions

DP23 is unusually broad because linguistic participation touches entry, expression, interpretation, retrieval, deliberation, and response. This section states the mechanism families that make the property operational and shows how they fit together. Each family is elaborated in the domain sections that follow, which remain the authoritative treatment of their subject.

The organizing claim is that language must be handled as infrastructure with declared guarantees rather than as a presentation-layer convenience. Six structural conditions follow from that claim.

5.1 Linguistic access layer

Participation begins with the ability to enter a space at all. DP23 requires that interface language, contribution language, and reading language be independently selectable; that script and Unicode integrity be preserved without normalization damage; that automatic language detection be overridable by the participant; and that non-text modalities including speech and sign be first-class contribution paths rather than accommodations.

Access is an entry condition, not a sufficiency condition. It establishes that language does not bar the door; it does not establish that participation will be understood or acted upon.

Elaborated in: *Linguistic Access as an Entry Condition*.

5.2 Original-canonical representation with translation provenance

Every derived representation is subordinate to its source. Originals are stored, displayed, and versioned; translations are paired views rather than replacements; and each derived artifact carries provenance stating source revision, target language, method, provider or model, timestamp, glossary version in force, and known uncertainty.

This is the mechanism that distinguishes DP23 from ordinary localization. A system that discards or hides originals has not implemented a weaker form of DP23; it has implemented the failure mode DP23 names.

Elaborated in: *Original-Language Integrity and Translation Provenance*.

5.3 Community-governed terminology

Meaning is issue-specific and contested. DP23 therefore requires glossaries as governed objects: versioned, attributable, forkable, and capable of holding plural or disputed entries rather than resolving them by fiat. Glossary governance has named stewards with bounded authority, and glossary changes leave a history because they retroactively affect how past contributions read.

Terminology control is a form of power. Making it explicit is what keeps it accountable.

Elaborated in: *Issue-Specific Multilingual Glossaries*.

5.4 Contestability and linguistic due process

Any interpretive act that affects a participant must be challengeable by that participant in a language they can use. This includes translations, cultural framing, classifications, cluster assignments, and summaries. Challenges have owners, timelines, evidence requirements, visible resolution history, and escalation paths, and notices of enforcement or appeal are themselves available in usable languages.

Due process that exists only in a dominant language is not due process for the people most likely to need it.

Elaborated in: *Community Stewardship, Contestability, and Due Process*.

5.5 Cross-language semantic interoperability

Shared understanding requires that meaning move across languages without routing all reasoning through a single dominant language. DP23 therefore requires multilingual search and retrieval, language-plural semantic representations, cross-language clustering that preserves minority positions, and monitoring for semantic drift between languages over time.

Where semantic equivalence cannot be established, the gap is surfaced rather than smoothed over (DP7, DP22).

Elaborated in: *Cross-Language Semantic Interoperability*.

5.6 Traceable sensemaking with bounded AI authority

Participation must connect to outcome. DP23 requires contribution-level traces, influence labels, provenance for recommendations and decisions, and identifiable institutional responsibility for response. Multilingual AI operates as advisory infrastructure: it may translate, cluster, and summarize, but it does not ratify consensus, determine representation, or issue final interpretive authority. Human ratification is required for high-impact multilingual outputs, and absence of participation is reported rather than concealed.

Elaborated in: *Multilingual AI as Advisory Infrastructure, Traceable Collective Sensemaking, and Representation, Missing Constituencies, and Epistemic Balance*.

5.7 Zone-scoped participation boundaries

Global reach is not the same as global authority. DP23 scopes participation to issues, with zone types that declare who may contribute, in what languages, under what stewardship, with what translation obligations, and with what relationship to decision authority. Cross-zone semantic compatibility allows zones to reference one another without merging their governance.

Elaborated in: *Global Issue Participation and Sociotechnical Zones for Multilingual Participation*.

5.8 Structural conditions across all mechanisms

Four conditions apply to every mechanism above and should be read as constraints rather than features:

- **Non-substitution.** No derived representation may present itself as the source.
- **Declared uncertainty.** Confidence and limitation travel with artifacts across languages, modalities, and systems.
- **Per-language accountability.** Quality, review capacity, and error rates are reported by language, because aggregate metrics conceal exactly the failures DP23 targets.
- **Bounded authority.** Every interpretive role, human or machine, has stated scope, term, and challenge path.

Example: A contribution submitted in Quechua is stored as canonical, paired with a machine translation labeled with model, glossary version, and low-confidence markers on two contested

terms. A steward annotates cultural context. The contribution is clustered with related positions across four languages, cited in a synthesis with a visible trace back to the original, and the participant sees that their point was incorporated into a recommendation and which institution is responsible for responding.

What this feels like: You speak as yourself, and the system carries your meaning forward with its uncertainty intact.

Without this: Multilingual interfaces produce the appearance of global participation while a single language quietly determines what was understood.

6. Linguistic Access as an Entry Condition

The Meta-Layer should provide the technical and governance conditions required for people to enter an issue space without being excluded by language.

Linguistic access is an entry condition.

It is not a guarantee that participation will be equitable, understood correctly, or acted upon.

Those outcomes require downstream semantic integrity, community oversight, representation analysis, and feedback.

6.1 Interface Language, Contribution Language, and Reading Language

DP23 distinguishes at least three language settings:

- **Interface language:** The language used for navigation, controls, instructions, consent, and governance information
- **Contribution language:** The language or languages used by the participant to express their input
- **Reading language:** The language or languages in which the participant prefers to review contributions and synthesis

These settings must not be conflated.

A participant may:

- navigate in one language
- contribute in another
- read translations in a third
- review originals in multiple languages

The system should support this plurality without requiring a single global language preference.

6.2 Unicode and Script Integrity

The Meta-Layer must preserve:

- writing systems
- diacritics
- bidirectional text
- punctuation
- line structure
- non-Latin scripts
- mixed-language contributions
- transliteration where participant-selected
- names and terms that should not be translated

Text normalization must not silently alter identity, meaning, or authorship.

6.3 Language Detection with Participant Override

Automated language detection may support usability, indexing, and translation.

It must not be treated as unquestionable.

Participants should be able to:

- confirm detected language
- correct detection
- identify mixed-language content
- specify dialect or regional form where relevant
- indicate that a contribution should not be translated automatically
- identify words or names that must remain in the original form

Low-confidence detection must not block participation.

6.4 Multimodal Linguistic Participation

DP23 operates in conjunction with DP21.

Participation should not be limited to typed text.

Where feasible, issue spaces should support:

- speech
- transcription
- sign-language video
- image-based contribution

- captioning
- audio description
- transliteration
- screen-reader compatibility
- low-bandwidth participation
- asynchronous voice contribution
- community-assisted transcription

Multimodality is especially important where literacy, disability, script support, oral tradition, device access, or connectivity would otherwise exclude participation.

7. Original-Language Integrity and Translation Provenance

Original contributions are the canonical civic record.

Every derived representation must remain connected to:

- the original contribution
 - the original language
 - the relevant revision
 - the actor who submitted it
 - the translation system or reviewer
 - the model and version used
 - the glossary version applied
 - uncertainty notes
 - corrections and challenges
 - creation and modification timestamps
-

7.1 Original and Translation as Paired Views

Participants should be able to view:

- original only
- translation only, with clear labeling
- original and translation side by side
- translation history
- alternative translations
- cultural annotations

- contested terminology

A translation should never visually masquerade as the original.

7.2 Translation Confidence and Uncertainty

Where confidence measures are available, they should be presented carefully and not treated as mathematical truth.

More importantly, systems must expose qualitative uncertainty such as:

- ambiguous term
- missing cultural equivalent
- uncertain speaker intent
- contested political terminology
- idiomatic phrase
- region-specific meaning
- multiple plausible translations
- glossary conflict
- untranslatable or intentionally untranslated term

Uncertainty is not a system defect to be hidden.

It is contextual information required for trustworthy interpretation.

7.3 Translation Versioning

Translations must be versioned.

A new translation should be created when:

- the original contribution changes
- the glossary changes materially
- a challenge is accepted
- a model or translation provider changes
- a human-reviewed version supersedes an automated version
- cultural context changes the preferred interpretation

Historical versions should remain inspectable according to the issue's governance and privacy rules.

7.4 Provenance Across Derived Representations

Provenance must continue beyond translation.

A participant should be able to navigate from:

Original contribution

↓

Translation

↓

Classification

↓

Semantic cluster

↓

Summary statement

↓

Recommendation

↓

Decision

↓

Assigned responsibility

↓

Response or action

Each transition must be attributable, inspectable, and challengeable.

8. Issue-Specific Multilingual Glossaries

General-purpose translation is insufficient for many civic, scientific, cultural, political, legal, ecological, technical, and spiritual issues.

DP23 therefore establishes issue-specific multilingual glossaries as shared semantic infrastructure.

8.1 Purpose of an Issue Glossary

An issue glossary helps participants and AI systems understand:

- how key concepts are defined
- which translations are preferred
- which translations are acceptable
- which terms are contested
- which terms are deprecated
- which concepts lack direct equivalents
- how meaning varies by region or community

- when terms should remain untranslated
- what historical or political context affects interpretation

A glossary does not impose universal linguistic conformity.

It makes semantic choices visible and governable.

8.2 Plural and Contested Terminology

A glossary may contain multiple translations for the same concept.

Terms may be marked as:

- preferred
- acceptable
- contested
- community-specific
- institutionally used
- historical
- deprecated
- harmful in certain contexts
- lacking an equivalent
- intentionally untranslated

No single institution should be assumed to possess final authority over meaning.

8.3 Glossary Governance

Communities should be able to:

- propose terms
- propose translations
- annotate cultural nuance
- provide examples
- challenge preferred terminology
- identify harmful framing
- document minority usage
- review AI-generated suggestions
- approve or reject changes
- fork glossaries when genuine semantic disagreement cannot be reconciled

Glossary governance should identify:

- who may propose changes
 - who may review changes
 - who may approve changes
 - how conflicts are resolved
 - how affected communities are represented
 - when external expertise is appropriate
 - how emergency corrections are made
 - how prior versions remain accessible
-

8.4 Glossary Capture as a Failure Mode

Glossaries can become instruments of power.

A glossary may be captured when:

- an institution defines the only valid terminology
- dominant-language speakers approve all translations
- political terminology is normalized without disclosure
- an affected community is excluded from review
- a preferred translation suppresses disagreement
- AI-generated terms are accepted without competent review

DP23 requires glossary authority, conflicts of interest, revision history, and contested status to be visible.

9. Community Stewardship, Contestability, and Due Process

Multilingual systems cannot be governed only by model providers, developers, or professional translators.

Affected communities must have meaningful mechanisms to participate in interpretation.

DP23 treats contestability as core linguistic infrastructure.

Participants must be able to:

- challenge translations
- annotate cultural nuance
- correct classifications
- identify missing constituencies
- contest summaries

- audit the process
-

9.1 Translation Challenges

A participant may challenge a translation on grounds including:

- incorrect meaning
- omitted meaning
- added meaning
- tone distortion
- cultural distortion
- political bias
- harmful terminology
- incorrect identity terminology
- false certainty
- glossary misuse
- mistranslated name or place
- failure to preserve ambiguity

Challenges should allow:

- explanation
 - proposed correction
 - cultural or regional context
 - supporting sources
 - alternative translations
 - affected-community testimony
-

9.2 Cultural Annotations

Not every interpretive problem can be solved by replacing words.

Participants should be able to attach contextual annotations such as:

- “This phrase is respectful in one region but insulting in another.”
- “This term carries a historical meaning not preserved in the translation.”
- “There is no direct equivalent in the target language.”
- “The speaker is using irony.”
- “This term refers to a relationship rather than an individual category.”
- “This concept should remain untranslated.”

Annotations should remain distinguishable from corrections.

9.3 Classification Challenges

Participants should be able to contest how a contribution has been classified, including whether it is treated as:

- evidence
- opinion
- testimony
- proposal
- concern
- question
- objection
- minority view
- harmful content
- irrelevant content
- duplicate content
- consensus-supporting content

Classification changes must retain an audit trail.

9.4 Summary Challenges

Participants should be able to challenge summaries that:

- misrepresent their contribution
- omit a significant perspective
- falsely imply consensus
- merge incompatible positions
- rely excessively on one language
- ignore relevant cultural context
- overrepresent repeated or coordinated submissions
- fail to identify missing constituencies
- mistranslate cited evidence
- obscure uncertainty

Significant challenges should receive human or community review.

9.5 Linguistic Due Process

Decisions affecting participation, visibility, classification, reputation, or influence must be understandable in a language the affected participant can meaningfully use.

Appeal mechanisms must not exist only in the platform's dominant language.

Due process includes:

- understandable notice
 - access to relevant evidence
 - explanation of automated involvement
 - ability to respond in the participant's language
 - competent review
 - visible resolution
 - correction of downstream effects where feasible
-

10. Cross-Language Semantic Interoperability

DP23 requires the Meta-Layer to support discovery and reasoning across languages.

A contribution should not become invisible merely because it uses a different language from the search query, moderator, facilitator, or summarization system.

10.1 Search Across Languages

Participants should be able to:

- search in one language
- retrieve relevant contributions in other languages
- inspect why a result was considered relevant
- request translation on demand
- open the original
- filter by source language
- identify machine-expanded search terms
- challenge irrelevant semantic matches

Search should not require every contribution to be translated in advance.

10.2 Language-Plural Semantic Representations

The long-term architecture should support semantic representations capable of relating contributions across languages without treating one dominant language as the permanent internal truth layer.

English or another language may be used as an early operational bridge.

It must not become an irreversible architectural dependency.

Language-plural reasoning should preserve:

- concepts
 - relationships
 - uncertainty
 - contradictions
 - culturally specific categories
 - source-language provenance
 - alternative interpretations
-

10.3 Semantic Drift Across Languages

Concepts may appear equivalent while functioning differently across communities.

Systems must guard against semantic drift in which:

- the same translated term carries different assumptions
- legal categories are treated as internationally uniform
- political identities are mapped incorrectly
- local ecological knowledge is forced into scientific taxonomies
- spiritual or kinship concepts are reduced to Western individualist categories
- institutional vocabulary overrides lived terminology

Where semantic alignment is uncertain, the system should preserve distinction rather than manufacture equivalence.

11. Global Issue Participation

Some issues affect people across borders, languages, jurisdictions, ecosystems, and generations.

DP23 establishes the conditions under which an issue may become a shared global participation space.

11.1 Definition

Global issue participation is the capacity for people affected by, knowledgeable about, or meaningfully concerned with an issue to:

- contribute in languages they can meaningfully use
- understand relevant contributions in languages they can meaningfully understand
- examine the evidence and positions in the corpus
- identify missing perspectives
- participate in interpretation and synthesis
- challenge AI-generated representations
- propose responses
- trace institutional responsibility
- follow subsequent decisions and action

Global issue participation does not require universal agreement.

It provides shared infrastructure for informed plurality.

11.2 Issue Scope and Participation Boundaries

Not every issue is global.

Issue stewards should define:

- who is affected
- who has relevant knowledge
- which jurisdictions are implicated
- which communities require representation
- which languages require active support
- whether participation is open, invited, delegated, sampled, or hybrid
- how local and global authority interact
- what decisions the process can influence
- what decisions remain outside its scope

These boundaries must be visible and contestable.

11.3 Participation Is Not Representation by Default

A person contributing in a language does not automatically represent all speakers of that language.

A participant should not be presented as representing:

- a country
- an ethnic group
- a linguistic community
- an Indigenous nation
- a profession
- an age group
- a religious group
- an affected population

unless an explicit and legitimate representative relationship exists.

DP23 distinguishes individual voice, community membership, expertise, testimony, delegation, and formal representation.

11.4 Global Ambassadors and Language Stewards

The Meta-Layer may support distributed networks of:

- language stewards
- cultural stewards
- issue ambassadors
- regional facilitators
- community translators
- glossary reviewers
- accessibility contributors
- constituency liaisons

These roles should not be reduced to promotion.

Their civic responsibilities may include:

- explaining an issue locally
- recruiting underrepresented participants
- improving terminology
- reviewing translations
- identifying cultural distortion
- surfacing regional priorities
- connecting local discussions to the global corpus
- communicating results back to communities

- identifying participation barriers
- monitoring whether institutions respond

Steward authority must remain scoped, transparent, accountable, and revocable.

12. Multilingual AI as Advisory Infrastructure

DP23 anticipates extensive use of AI for:

- translation
- transcription
- semantic search
- clustering
- classification
- summarization
- question detection
- constituency-gap detection
- recommendation synthesis
- trace generation
- moderation support

AI may support multilingual participation.

It may not become an invisible authority over meaning.

12.1 In-Language Evaluation

Multilingual systems must not be evaluated only by:

- translating output back into English
- measuring surface fluency
- using dominant-language reviewers
- applying one universal benchmark
- checking whether the output “sounds natural”

Evaluation should include, where feasible:

- fluent or native reviewers
- affected-community review
- issue-specific terminology
- regional language variation

- ambiguity tests
 - cultural-context tests
 - political and historical terminology
 - preservation of tone and uncertainty
 - low-resource language performance
 - downstream classification and summary effects
-

12.2 Contextual Reasoning

AI systems should reason with relevant context, including:

- issue glossary
- community terminology
- source language
- contribution history
- cultural annotations
- speaker-selected identity terms
- relevant jurisdiction
- corpus disagreements
- known ambiguity

Context must not be invented when absent.

12.3 Uncertainty Handling

AI should be able to state:

- “Multiple translations are plausible.”
- “This concept has no direct equivalent.”
- “The interpretation depends on regional context.”
- “The source phrase may be ironic.”
- “The glossary contains a contested term.”
- “The classification confidence is low.”
- “The corpus does not contain enough participation from this constituency.”
- “This summary should not be treated as representative.”

Uncertainty must not automatically suppress content.

It should trigger review, annotation, or cautious presentation.

12.4 Community Participation

Communities should participate in:

- evaluation-set creation
- glossary development
- translation review
- classification review
- summary review
- error reporting
- model comparison
- harm assessment
- release decisions
- ongoing monitoring

AI developers alone cannot determine culturally appropriate system behavior.

12.5 Human Ratification

AI systems may propose:

- translations
- classifications
- clusters
- summaries
- participation gaps
- recommendations
- terminology

For high-impact issue processes, humans or legitimate community governance bodies should ratify material outputs before they are treated as authoritative.

AI proposes.

Communities interpret, challenge, and decide.

13. Traceable Collective Sensemaking

Participation becomes meaningful when contributors can see how their input moved through the process.

DP23 requires visible participation traces.

A contributor should be able to determine:

- where their input went
 - how it was grouped
 - what was learned
 - which recommendations changed
 - who is responsible for responding
 - what happens next
-

13.1 Contribution-Level Trace

For each contribution, the system should show, where applicable:

- original submission
 - translations
 - glossary terms applied
 - classifications
 - clusters
 - summaries citing or relating to it
 - challenges
 - recommendation links
 - decision links
 - action links
 - documented institutional responses
-

13.2 Influence Labels

Systems must not overstate causal influence.

Permitted labels may include:

- included in corpus
- translated
- classified as
- grouped under
- directly cited
- supporting evidence
- contradicting evidence
- minority perspective

- considered during synthesis
- contributed to revision
- referenced in recommendation
- no documented influence yet

The system should distinguish direct citation from semantic similarity.

13.3 Recommendation and Decision Traceability

Recommendations should identify:

- source clusters
- supporting contributions
- contradicting contributions
- languages represented
- missing constituencies
- challenge status
- revision history
- responsible party
- response deadline where applicable
- implementation status

Changes should explain:

- what changed
 - why it changed
 - which evidence or participation affected the revision
 - who authorized the change
-

13.4 Responsibility and Response

Participation processes should identify:

- who is responsible for reviewing input
- who may make decisions
- who must respond
- whether response is mandatory
- expected timing
- escalation pathways
- what happens if no action is taken

A participation system without responsibility mapping risks becoming performative consultation.

14. Representation, Missing Constituencies, and Epistemic Balance

The existence of a multilingual corpus does not establish representativeness.

DP23 requires systems to distinguish between:

- contribution volume
 - language diversity
 - geographic diversity
 - constituency diversity
 - affected-community participation
 - expertise diversity
 - representativeness
 - legitimacy
-

14.1 Missing Constituency Detection

Participants and facilitators should be able to identify missing constituencies, including:

- affected communities
- language communities
- geographic regions
- future generations
- professional groups
- local knowledge holders
- Indigenous peoples
- youth
- elders
- displaced populations
- people with disabilities
- low-connectivity communities
- nonparticipants directly affected by outcomes

AI may help surface likely gaps.

It must not infer sensitive identity or claim complete constituency knowledge without appropriate evidence and consent.

14.2 Visibility Without Tokenization

Minority and underrepresented perspectives should remain visible without being reduced to symbolic inclusion.

Systems should avoid:

- presenting one participant as the voice of a community
 - creating a single “minority view”
 - ranking perspectives solely by volume
 - using language count as a proxy for inclusion
 - treating translation availability as evidence of participation
 - displaying underrepresented views without protecting vulnerable contributors
-

14.3 Weighting and Aggregation

DP23 does not prescribe one global method for weighting contributions.

Issue communities may choose among:

- one-person-one-contribution approaches
- deliberative sampling
- constituency weighting
- expertise weighting
- affectedness weighting
- random selection
- delegated participation
- reputation-informed review
- unweighted open corpora
- hybrid approaches

Any weighting system must make its assumptions, authority, and consequences visible.

15. Sociotechnical Zones for Multilingual Participation

Multilingual participation should be configurable through sociotechnical zones.

A zone may define:

- supported interface languages
- supported contribution modes
- translation providers

- glossary governance
 - steward roles
 - confidentiality conditions
 - review requirements
 - AI usage
 - publication rules
 - representation requirements
 - response obligations
 - data residency
 - cultural protocols
 - jurisdictional requirements
-

15.1 Example Zone Types

Open Multilingual Zone

- Public participation
- Machine translation on demand
- Community challenges
- Public corpus
- Lightweight glossary governance

Affected-Community Zone

- Participation prioritized for directly affected groups
- Community-approved terminology
- Strong cultural protocols
- Controlled data access
- Human-reviewed synthesis

High-Stakes Deliberation Zone

- Verified participation roles
- Professional interpretation
- AI-assisted but human-ratified summaries
- Formal appeals
- Full provenance and audit
- Institutional response obligations

Indigenous Knowledge Zone

- Community-defined access
- Restrictions on translation or reuse
- Community-controlled terminology
- Traditional knowledge protections
- Explicit consent for AI processing
- Local governance authority

Youth Participation Zone

- Age-appropriate interfaces
- safeguarding
- accessible language
- guardian or institutional protocols where necessary
- youth governance roles
- protection from extractive data use

15.2 Cross-Zone Semantic Compatibility

A term, identity, reputation signal, or governance decision may not carry the same meaning across zones.

Systems must signal when:

- glossary definitions differ
- translation policies differ
- participation standards differ
- confidentiality conditions differ
- AI use differs
- representative authority differs
- cultural restrictions apply

Cross-zone interoperability requires semantic signaling, not forced uniformity.

16. Governance, Accountability, and Agency Surfaces

Multilingual infrastructure distributes interpretive authority whether or not it declares that it is doing so. Someone decides which languages are supported, which glossary entry prevails, which steward is trusted, which summary is published, and which institution must answer. DP23 requires those decisions to have named surfaces and reachable levers.

16.1 Accountability surfaces

- **A language support statement** per zone or issue, listing supported contribution and reading languages, translation methods in use, review capacity available, and known quality limits by language
- **A translation provenance record** attached to every derived representation, sufficient to reproduce how it was produced and by what version of what system
- **A glossary change history** identifying proposer, approver, prior term, rationale, and effective date, including contested entries left unresolved
- **A steward register** naming language stewards, ambassadors, and reviewers, with scope, term, appointment route, and removal procedure
- **A challenge docket** showing open and resolved translation, classification, and summary challenges, with timelines and outcomes reported per language
- **An AI disclosure record** covering models, prompts where disclosable, sampling and evaluation practice, and error rates broken out by language
- **A representation statement** accompanying any claim of collective input, declaring which constituencies are present, which are absent, and what the absence means for the claim
- **An institutional response record** identifying who received a synthesis, what they were asked, what they answered, and by when

16.2 Participant agency surfaces

Accountability is only meaningful where participants can act. DP23 expects the following powers to be available in a language the participant can actually use:

- choose contribution, reading, and interface languages independently, and change them without losing history
- see the original alongside any translation of their own or another's contribution
- challenge a translation, cultural framing, classification, cluster placement, or summary that concerns them, and see the resolution
- propose glossary terms, register a dissenting definition, or mark a term as contested
- decline translation, indexing, or summarization of contributions that carry cultural, safety, or sovereignty constraints
- trace their own influence: whether a contribution was read, grouped, cited, incorporated, or left without documented effect
- report that a constituency is missing, and have that report enter the representation record
- appeal to a body that did not make the original interpretive decision

- exit or fork a glossary, zone, or stewardship arrangement with continuity of their own contributions (DP7, DP20)

Example: A participant contributing in Amharic receives notice that a summary characterized their position as opposition to a policy when they had raised a procedural objection. They file a summary challenge in Amharic, a steward and a second reviewer assess it, the summary is annotated pending resolution, and the correction is recorded with its reasoning.

What this feels like: Interpretation is something you can reach and argue with, in your own language.

Without this: Participants are translated, classified, and summarized by a process they can neither see nor answer.

Failure mode: **due process in translation only**, where remedies formally exist but are reachable exclusively through a dominant language.

17. Incentives and Power Analysis

Linguistic infrastructure is expensive, unevenly funded, and commercially concentrated. DP23 assumes that the economics of language technology will shape participation unless the incentives are designed against.

17.1 Where power concentrates

- **Model and translation providers.** A small number of vendors determine quality, cost, and availability for most language pairs, and their roadmap decisions become participation boundaries.
- **High-resource language communities.** Better tooling produces better representation, which produces more data, which produces better tooling. The advantage compounds without anyone intending it.
- **Glossary and terminology stewards.** Whoever defines contested terms shapes what can be argued.
- **Synthesis producers.** Those who write or generate summaries determine what institutions read.
- **Institutions receiving input.** Where response is discretionary, the power to ignore is the decisive power.
- **Platform operators.** Retrieval and ranking determine which language communities are practically visible (DP22).

17.2 Predictable incentive distortions

- **Coverage over quality:** claiming many languages is cheaper and more marketable than serving a few languages well, and aggregate metrics reward the claim.

- **Low-resource neglect:** languages with small commercial value receive unusable quality, and the disclosure of that gap is itself deprioritized.
- **Unpaid stewardship:** communities are asked to correct machine output for free, effectively subsidizing vendors with volunteer expertise.
- **Extraction of corpora:** community contributions, glossaries, and corrections become training data that improves systems the community does not control (DP4, DP20).
- **Legitimacy laundering:** institutions purchase multilingual participation to authorize decisions already made, because breadth of input is easier to demonstrate than responsiveness.
- **Steward capture:** recognition and authority attract actors whose interest is positional rather than linguistic.
- **Translation as cost center:** review capacity is cut first under budget pressure, and the failure surfaces only for the least powerful participants (DP17).

17.3 Counter-incentives DP23 expects

- fund per-language review capacity explicitly, and publish that capacity so shortfalls are visible before they cause harm
- compensate language stewards, translators, and cultural mediators as infrastructure roles rather than as community goodwill (DP9)
- require per-language quality reporting, so that coverage claims cannot conceal unusable service
- treat community-contributed glossaries, corrections, and corpora as community-owned assets with consent-bound reuse terms
- prioritize investment toward the languages with the largest gap between affected population and current quality, rather than toward the easiest gains
- bind institutional response commitments to participation requests, so that soliciting input carries an obligation
- rotate and bound steward authority, and pair recognition with accountability
- measure success by whether influence traces reach participants, not by volume of translated content

Example: A civic body funding a multilingual consultation publishes its review capacity per language, contracts paid stewards for the four lowest-resource languages involved, and commits in advance to a dated response naming the office responsible.

Failure mode: **participation as procurement**, where multilingual input is bought to satisfy a legitimacy requirement while the cost of understanding it is never funded.

18. Community Signals Informing DP23

DP23 should remain open to refinement through multilingual community input.

Early signals to validate include:

- People want to contribute in the language in which they can express themselves most precisely.
- Participants want access to translations without losing access to original wording.
- Communities do not want AI systems to silently decide culturally sensitive terminology.
- Participants want mechanisms to correct translation and classification errors.
- Low-resource language communities want involvement in evaluation, not merely access to finished products.
- Contributors want to know whether their input affected a summary, recommendation, or decision.
- Communities want missing voices identified rather than hidden behind claims of global participation.
- Institutions need multilingual synthesis, but participants require traceability and meaningful response.
- Language stewards want recognition, authority boundaries, support, and accountability.
- Participants want systems to preserve disagreement rather than manufacture consensus.

These signals should be tested through workshops, pilots, calls for input, and language-specific consultations.

DP23 should not be declared complete based only on English-language deliberation about multilingual inclusion.

19. Evaluation Criteria

DP23 requires continuous evaluation across languages and issue contexts.

Evaluation should measure:

- meaning preservation
- omission and addition rates
- terminology accuracy
- tone preservation
- uncertainty preservation

- cultural-context integrity
- quality differences between languages
- challenge rates
- correction acceptance rates
- review times
- classification disparities
- search visibility across languages
- summary citation rates by language
- minority-view preservation
- representation gaps
- contribution trace completeness
- participant understanding
- institutional response rates
- participant perception of meaningful influence

Aggregate quality scores must not conceal poor performance in specific languages.

Communities should be able to inspect:

- which languages perform well
- which languages perform poorly
- what review capacity exists
- what errors recur
- what changes have been made
- which systems or providers are responsible

Evaluation should also be reviewed for its own blind spots. Metrics produced only in a dominant language, assessed only by well-resourced reviewers, or reported only in aggregate will systematically miss the failures DP23 exists to detect. Evaluation design is therefore subject to the same contestability requirements as translation and classification.

20. Implementation Patterns

The following patterns recur across systems that implement DP23 successfully. They are offered as reusable approaches rather than as requirements, and they are distinct from the phased rollout described later in this draft.

20.1 Paired-view rendering

Display original and translation together by default, with the original visually primary. Where space is constrained, the original remains one interaction away and is never removed from the object.

20.2 Provenance envelopes on derived artifacts

Attach a structured provenance block to every translation, summary, cluster, and classification, carrying source revision, method, model or translator identity, glossary version, timestamp, and confidence. Treat a missing envelope as a defect, not as a default.

20.3 Confidence surfacing rather than confidence hiding

Render uncertainty at the point of reading: inline markers on low-confidence spans, span-level rather than document-level scores where possible, and explicit notices where a term has no established equivalent.

20.4 Glossary as a versioned object

Store terminology as a governed artifact with identifiers, versions, plural entries, and change history. Bind each translation to the glossary version that produced it, so that later term changes do not silently rewrite past meaning.

20.5 Challenge-first interfaces

Place the challenge affordance next to the interpretive artifact rather than in a support menu, accept challenges in the participant's own language, and show challenge status on the artifact while review is pending.

20.6 Steward workflows with bounded authority

Route challenges to stewards with declared scope, require a second reviewer for culturally sensitive or high-impact decisions, publish outcomes, and rotate assignments to prevent positional entrenchment.

20.7 Language-partitioned dashboards

Report quality, latency, review capacity, challenge volume, and correction acceptance per language by default. Make aggregate views the exception rather than the presentation layer.

20.8 Gap detection as a first-class feature

Compare participating constituencies against those affected by an issue, and render the difference as part of any synthesis rather than as an appendix.

20.9 Influence receipts

Give each contributor a durable trace of what happened to their contribution: read, grouped, cited, incorporated, contested, or unused, with links to the artifacts involved (DP18, DP22).

20.10 Human ratification gates

Require named human approval before a multilingual synthesis is published as representing collective input, and record who ratified what on which evidence.

20.11 Graceful refusal

Where a community has restricted translation, indexing, or summarization, render an explicit notice of restriction rather than an empty result, so that absence is legible rather than mistaken for silence.

21. Foresight and Failure Design

DP23 treats multilingual inclusion as an adversarial and governance-sensitive design problem.

This section applies minefield thinking: naming the ways language infrastructure fails before those failures acquire constituencies.

Systems should anticipate how language infrastructure may be captured, manipulated, degraded, or used to create false legitimacy.

21.1 Governance Pre-Mortems

Communities should periodically ask:

- How could dominant-language actors capture the process?
 - How could AI-generated multilingual submissions overwhelm human voices?
 - Which languages are receiving lower-quality service?
 - How could translation challenges be brigaded?
 - How could glossary terms be politically manipulated?
 - How could cultural stewards become unaccountable gatekeepers?
 - How could institutions claim global legitimacy from a narrow participant pool?
 - How could summaries conceal unresolved disagreement?
 - How could a vulnerable community's knowledge be extracted or exposed?
 - How could participation traces reveal sensitive information?
 - How could a model-provider change alter historical interpretation?
 - What happens when a translation provider becomes unavailable?
-

21.2 Preventative Friction

Appropriate safeguards may include:

- rate limits
- proof-of-humanity for influence-sensitive actions
- steward review
- glossary-change waiting periods
- constituency review
- translation quality thresholds
- human approval for high-stakes summaries
- disclosure of coordinated campaigns
- restrictions on bulk AI-generated submissions
- culturally governed access controls
- delayed publication
- consent renewal
- provenance requirements

Preventative friction should be proportional to risk.

21.3 Exit, Fork, and Refusal

Communities and participants should be able to:

- decline machine translation
- decline AI summarization
- restrict reuse
- withdraw from an active process where governance permits
- preserve an original while retracting a public translation
- fork a glossary
- fork a summary
- establish a separate cultural interpretation
- leave a zone without losing access to their contribution history
- request that certain knowledge remain local or untranslated

Refusal is part of linguistic agency.

21.4 Emergency Mechanisms

High-impact systems should support mechanisms to:

- pause faulty translations
 - withdraw a harmful glossary version
 - suspend an unreliable model
 - stop publication of a misleading summary
 - flag a corpus as nonrepresentative
 - freeze downstream recommendations
 - notify affected contributors
 - restore prior validated versions
 - initiate independent review
-

22. Relationship to Other Desirable Properties

DP23 is a distinct property, but it depends on and strengthens other DPs.

DP1 – Federated Authentication & Accountability

Ensures participants, language stewards, AI systems, and institutions can be accountable without requiring universal real-name identity.

DP2 – Participant Agency and Empowerment

Establishes language choice, correction, refusal, contestation, and influence traceability as expressions of agency.

DP3 – Adaptive Governance

Supports multilingual governance participation and adaptation as communities, languages, and issue contexts evolve.

DP4 – Data Sovereignty and Privacy

Protects contributions, translations, cultural knowledge, identity information, and consent across jurisdictions and language services.

DP7 – Simplicity and Interoperability

Provides technical and semantic interoperability across languages, interfaces, tools, and communities.

DP8 – Collaborative Environment and Meta-Communities

Enables meta-communities to collaborate across linguistic boundaries while maintaining local identity and governance.

DP10 – Education

Supports learning, explanation, literacy, and access to knowledge across languages and modalities.

DP11 – Safe and Ethical AI

Requires multilingual AI to remain transparent, evaluated in-language, context-aware, uncertain where appropriate, and accountable.

DP12 – Community-Based AI Governance

Provides the governance basis for translation challenges, cultural annotations, classification corrections, glossary stewardship, and summary review.

DP13 – AI Containment

Constrains multilingual AI to approved roles, data boundaries, models, glossaries, and governance rules.

DP14 – Trust and Transparency

Makes language mediation, AI involvement, uncertainty, governance authority, and interpretive decisions visible.

DP15 – Security and Provenance

Preserves original-language records and the lineage of translations, classifications, summaries, recommendations, and actions.

DP18 – Feedback Loops and Reputation

Supports correction loops, steward reputation, translation improvement, and visible institutional response.

DP20 – Community Ownership

Supports community ownership of glossaries, language resources, semantic models, cultural annotations, and issue corpora.

DP21 – Multi-Modal

Extends participation beyond written text through voice, sign, image, transcription, and accessible interfaces.

DP22 – Civic Memory & Epistemic Continuity

Preserves multilingual civic memory across time without collapsing original expressions into one normalized language.

DP23 should not be absorbed into any single one of these properties.

Its unique role is to establish **universal linguistic participation as a condition of legitimate collective sensemaking**.

23. Non-Goals and Explicit Boundaries

DP23 does not attempt to:

Create one universal language

The goal is interoperability without linguistic homogenization.

Guarantee perfect translation

No translation system can fully preserve every cultural, emotional, historical, and contextual dimension.

DP23 requires uncertainty, correction, and provenance rather than claims of perfection.

Force every contribution to be translated

Communities and participants may restrict translation for privacy, safety, cultural, legal, or sovereignty reasons.

Treat AI as the final authority on meaning

AI supports interpretation.

It does not possess exclusive authority to determine meaning, representation, consensus, or legitimacy.

Eliminate all linguistic disagreement

Some disagreement is irreducible and should remain visible.

Make every issue globally open

Participation scope remains issue-specific and governed.

Claim global representativeness from language diversity alone

A corpus containing many languages may still exclude affected constituencies or reproduce unequal power.

Replace professional interpreters, translators, facilitators, or cultural mediators

AI may augment these roles but should not erase the need for human competence and accountable judgment.

Extract protected or sacred knowledge

Communities retain authority over whether knowledge may be translated, indexed, summarized, or shared.

Impose one global glossary

Terminology may remain plural, contested, local, and forkable.

Equate participation with decision authority

A person's ability to contribute does not automatically determine the legal or institutional authority of the resulting process.

By naming these boundaries explicitly, DP23 remains ambitious without becoming universalizing or extractive.

24. Minimum DP23 Alignment (Non-Normative)

Minimum alignment is not a checklist of translation features.

It is the threshold at which a system can reasonably claim to support multilingual participation without making original expression subordinate to machine-generated interpretation.

A minimally aligned system should:

Preserve Originals

- Store and display original contributions
- Preserve source language and revision history
- Never silently replace originals with translations

Support Participant Language Choice

- Allow contribution in multiple languages
- Allow participant correction of language detection
- Separate interface, contribution, and reading language

Provide Translation Provenance

- Identify machine-generated translations
- Record source revision, target language, provider, model, and timestamp

- Display uncertainty or limitations where known

Enable Contestability

- Allow participants to challenge translations
- Allow cultural annotations
- Allow correction of material classification errors
- Maintain visible resolution history

Support Issue Context

- Allow issue-specific terminology or glossary support
- Record which glossary version affected a translation
- Permit contested terminology

Preserve Traceability

- Link summaries and classifications to source contributions
- Avoid uncited published synthesis
- Provide at least a basic indication of where contributions were used

Avoid False Claims of Representation

- Do not describe a multilingual corpus as globally representative without evidence
- Identify important participation gaps where known
- Distinguish individual contribution from community representation

Maintain Human Oversight

- Require competent review for high-impact multilingual outputs
- Ensure AI-generated interpretations remain challengeable
- Prevent autonomous AI from issuing final claims of consensus or representation

A system that offers automatic translation but does not preserve originals, expose provenance, permit correction, or show downstream influence may be multilingual in appearance, but it is not aligned with DP23.

25. Strong DP23 Alignment

A strongly aligned implementation would additionally provide:

- multilingual semantic search
- language-plural embeddings or equivalent semantic infrastructure

- issue-specific multilingual glossaries
 - community language stewards
 - cultural annotation systems
 - translation and summary versioning
 - multilingual due process
 - in-language quality evaluation
 - constituency-gap detection
 - minority-view preservation
 - cross-language clustering
 - traceable summaries
 - recommendation provenance
 - institutional responsibility mapping
 - participant-facing influence traces
 - multimodal participation
 - community control over sensitive knowledge
 - privacy-preserving translation workflows
 - auditable AI models and prompts
 - transparent correction metrics by language
 - explicit mechanisms for glossary and governance forks
-

26. Failure Conditions

A system should not claim DP23 alignment when:

- original contributions are discarded after translation
- all reasoning is invisibly routed through one dominant language
- translations are presented as authoritative originals
- participants cannot challenge material errors
- low-resource languages consistently receive unusable quality without disclosure
- summaries have no traceable sources
- AI-generated consensus is presented without representation analysis
- institutions collect input without identifying responsibility or next steps
- cultural knowledge is translated or reused against community rules
- language stewards have unbounded or invisible authority
- glossary changes have no history

- participants cannot understand enforcement or appeals in a usable language
 - multilingual participation is used to claim legitimacy while affected constituencies remain absent
-

27. Initial Implementation Path

DP23 may be developed in phases.

Phase 1 – Linguistic Access

- Unicode-safe contribution storage
- language detection with participant override
- reading-language preferences
- on-demand translation
- original and bilingual views
- translation provenance
- translation caching
- basic challenges

Phase 2 – Community Stewardship

- cultural annotations
- correction proposals
- language steward roles
- issue-specific multilingual glossaries
- contested terminology
- version history
- multilingual review workflows

Phase 3 – Cross-Language Sensemaking

- multilingual semantic search
- cross-language clustering
- evidence and position mapping
- traceable summaries
- minority-view preservation
- participation-gap analysis
- summary challenges

Phase 4 – Global Issue Participation

- constituency mapping
- global and regional ambassador networks
- structured issue participation zones
- multilingual deliberation
- recommendation provenance
- responsibility assignment
- response tracking
- contribution influence traces

Phase 5 – Federated Public Infrastructure

- interoperable glossary protocols
- portable linguistic preferences
- federated language stewardship
- shared evaluation resources
- multilingual civic-memory standards
- cross-platform issue corpora
- institution-facing response interfaces
- decentralized governance of language infrastructure

28. Proposed Constitutional Statement

No person's ability to participate in collective sensemaking should depend on the language they speak. The Meta-Layer must preserve original expression, enable understanding across languages, make machine interpretation visible and contestable, protect cultural and contextual integrity, and show how participation influences shared understanding and action.

29. Open Questions and Future Work

DP23 is deliberately ambitious, and several of its conditions do not yet have settled technical or governance answers. The questions below are open in the sense that reasonable implementers currently disagree, not merely unimplemented.

1. What minimum schema should express translation provenance so that it survives transfer between systems, modalities, and archives (DP7, DP21, DP22)?

2. How should span-level uncertainty be represented in a way that is machine-readable and readable by a participant who is not a linguist?
3. What obligations attach when a model provider changes versions, given that historical interpretations may shift retroactively?
4. How should glossaries be forked and reconciled, and what happens to translations bound to an abandoned glossary version?
5. Who holds authority over contested terminology when the affected community spans jurisdictions with conflicting norms?
6. How can language steward authority be bounded and rotated without losing the accumulated contextual expertise that makes stewardship valuable?
7. What evidentiary standard should a translation challenge meet, and how is that standard kept from becoming a barrier to the participants most affected?
8. How should quality thresholds be set for low-resource languages, and what is the correct behavior when quality falls below usable levels: disclosure, refusal, human-only routing, or something else?
9. How can constituency-gap detection identify missing voices without profiling participants or inferring sensitive attributes (DP4)?
10. What aggregation methods respect linguistic plurality without either flattening minority positions or granting them disproportionate weight?
11. How should influence traces be provided without exposing participants to retaliation in contexts where contribution carries risk?
12. What consent framework should govern reuse of community glossaries, corrections, and multilingual corpora as training data (DP20)?
13. How should sign languages, oral traditions, and non-written knowledge forms be represented as canonical originals rather than as transcriptions?
14. What is the appropriate relationship between issue-scoped participation and formal decision authority, and how is that boundary communicated across languages?
15. How can cross-language semantic drift be detected over long periods, and who is responsible for acting on it?
16. What accountability applies when an institution solicits multilingual participation and then does not respond?
17. How should emergency mechanisms operate when translation infrastructure fails during a time-critical civic process?
18. How can DP23 alignment be evaluated by the communities it concerns rather than by the systems claiming it?

These questions should be worked through with multilingual participation rather than resolved in a dominant language and translated afterward.

30. Path Toward ML-RFC

DP23 is currently an ML-Draft. It has an unusually broad scope, and its advancement is expected to proceed through narrow, testable components rather than as a single specification.

Advancement toward ML-RFC status SHOULD require:

- a minimal interoperable schema for translation and transformation provenance, implemented by at least two independent systems
- a glossary object format supporting versioning, plural entries, contested status, attribution, and forking
- a documented challenge protocol covering translation, classification, and summary challenges, with resolution records
- an influence trace format connecting contributions to summaries, recommendations, and decisions
- per-language quality and review-capacity reporting demonstrated in a live deployment
- a constituency-gap reporting method that operates without sensitive-attribute inference
- at least one multilingual deliberation pilot in which low-resource language participants took part in evaluation, not only in contribution
- evidence that human ratification gates functioned under time pressure rather than being bypassed

ML-RFC promotion SHOULD be contingent on:

- rough consensus reached through multilingual process, including participants from low-resource language communities
- demonstrated interoperation with DP14 transparency, DP15 provenance, and DP22 memory pathways
- documented failure cases, including at least one case where the system correctly refused to claim representativeness
- alignment with professional translation, interpretation, and cultural mediation practice rather than displacement of it
- evidence that participants could exercise challenge and trace rights in their own languages

Early ML-RFC candidates may focus on translation provenance and glossary objects, since both are narrowly scoped, independently useful, and verifiable without requiring the full participation architecture.

DP23 should not be promoted on the strength of a reference implementation whose evaluation was conducted only in English.

31. Closing Orientation

DP23 is not a requirement that everyone speak the same language, adopt the same categories, or reach the same conclusion.

It establishes the conditions under which people can inhabit the same civic and epistemic space while preserving linguistic and cultural difference.

The Meta-Layer should not achieve global participation by translating humanity into one voice.

It should enable humanity's many voices to become mutually legible, collectively navigable, and meaningfully connected to responsibility and action.
