

DP10 – Education

The Meta-Layer supports dynamic, AI-powered learning tools that adapt to you — not the other way around.

Purpose of This Draft

This draft articulates Desirable Property 10 (DP10) as the condition under which participants can learn, onboard, grow, teach, and co-create within the meta-layer without being excluded by technical complexity, jargon, institutional gatekeeping, or static training models.

DP10 defines education as a continuous civic function of the meta-layer. It includes onboarding, tool literacy, shared vocabulary, AI-assisted guidance, peer learning, formal and informal curricula, community knowledge-sharing, and portable recognition of learning through credentials such as PEARL digital badges.

The central claim is that the meta-layer cannot become public infrastructure unless people can understand it, use it, critique it, teach it, and evolve with it.

Education in the meta-layer is not only content delivery. It is contextual learning across the web. Participants learn by doing, annotating, bridging, reflecting, teaching, and contributing inside live environments.

If DP10 is weak, predictable failures follow: only technical insiders participate meaningfully; newcomers become dependent on intermediaries; communities fragment around misunderstood terminology; AI tools create passive users rather than capable participants; credentials fail to travel; and the meta-layer becomes infrastructure people inhabit without understanding.

DP10 connects directly to:

- DP2, participant agency and empowerment
- DP3, adaptive governance
- DP7, interoperability
- DP8, community-defined participation and governance zones
- DP9, developer and community incentives
- DP11, safe and ethical AI
- DP14, trust and transparency
- DP18, feedback loops and reputation
- DP19, amplifying presence and community engagement
- DP21, multi-modal interactions and experiences

DP10 does not prescribe one curriculum, credentialing system, or educational institution. It defines the minimum conditions under which learning remains accessible, adaptive, community-grounded, and portable.

Problem Statement

Today's web gives people access to information but does not reliably help them develop understanding.

Search results, feeds, tutorials, chats, videos, FAQs, and documentation are abundant, but learning remains fragmented. Participants are expected to navigate unfamiliar tools, opaque algorithms, AI-generated content, misinformation, privacy risks, governance procedures, and technical vocabularies without coherent guidance.

This produces recurring failures:

- onboarding assumes prior literacy
- documentation is written for insiders
- communities use the same words differently
- AI assistants answer questions without building durable understanding
- educational credentials remain trapped in institutions or platforms
- youth and families lack safe pathways into civic digital literacy
- informal learning is rarely recognized
- participants learn tools but not governance, ethics, or agency

DP10 reframes education as an embedded layer of the meta-layer. Learning must happen where participants already act, not only in separate courses or manuals.

A healthy DP10 implementation must answer:

- How does a new participant understand what the meta-layer is?
- How do they learn what tools, rights, risks, and responsibilities exist?
- How do communities develop shared vocabulary?
- How does AI assist learning without replacing participant judgment?
- How are skills recognized across systems?
- How do youth, families, educators, builders, municipalities, and communities learn differently?
- How does learning feed back into governance and reputation?

Without DP10, the meta-layer risks becoming technically powerful but socially illegible.

Threats and Failure Modes

3.1 Onboarding cliffs

New participants encounter too much complexity too quickly.

Example: A participant installs a meta-layer tool and is immediately asked to understand zones, overlays, credentials, bridges, governance rules, and AI agents without a guided path.

Why this matters: Complexity without scaffolding produces dependency, abandonment, or misuse.

3.2 Documentation for insiders

Guides assume technical, crypto, governance, or AI literacy.

Example: A glossary explains “semantic interoperability” by referencing other specialized terms, leaving non-technical participants further behind.

Why this matters: Public infrastructure must be learnable by the public.

3.3 Vocabulary fragmentation

Different communities use the same terms differently, or different terms for the same concept.

Example: “Bridge,” “zone,” “overlay,” “agent,” and “credential” mean different things across technical, civic, educational, and AI governance groups.

Why this matters: Shared action requires shared meaning. Misaligned language becomes a governance risk.

3.4 Passive AI tutoring

AI learning assistants provide answers without helping participants build durable understanding.

Example: A participant asks how to evaluate provenance. The assistant gives an answer, but does not teach the participant how to inspect sources next time.

Why this matters: Education should increase agency, not deepen reliance.

3.5 Credential theater

Badges or certificates are issued without meaningful evidence of learning, reflection, or contribution.

Example: A participant receives a badge for clicking through onboarding screens, but cannot explain core safety, governance, or agency concepts.

Why this matters: Credentials should signal capability, not participation theater.

3.6 Institutional lockout

Learning that happens in communities, homes, peer groups, or informal settings is not recognized by schools, employers, or civic institutions.

Example: A youth participant builds high-quality annotations and civic maps, but those skills cannot be translated into academic credit or professional recognition.

Why this matters: Lifelong learning must become portable and institutionally legible without being captured by institutions.

3.7 Youth capture

Educational pathways target young people for adoption rather than empowering them as critical participants and stewards.

Example: A school campaign introduces the meta-layer as a product to use, but not as a civic system students can question, shape, and govern.

Why this matters: Youth education must cultivate agency, not brand loyalty.

3.8 Family and community exclusion

Education is treated as school-only or technical-only, excluding parents, caregivers, elders, libraries, local groups, and informal networks.

Example: A parent wants to help their child navigate synthetic media but cannot find accessible materials outside a technical whitepaper.

Why this matters: Digital literacy is intergenerational and community-based.

3.9 Learning without governance connection

Participants learn how to use tools but not how to participate in rule-making, feedback, appeals, or stewardship.

Example: A tutorial teaches annotation but not how annotations are moderated, contested, or incorporated into shared knowledge.

Why this matters: Tool literacy without governance literacy leaves participants disempowered.

3.10 Over-gamification

Learning badges and points reward completion, speed, or volume instead of reflection and meaningful contribution.

Example: A learner races through modules to collect badges, while slower reflective participants appear less accomplished.

Why this matters: Learning systems must not reproduce attention-economy incentives.

3.11 Static curricula in dynamic environments

Educational materials become outdated as tools, policies, and risks change.

Example: A safety guide describes old AI disclosure patterns while newer agents operate with different permissions and risks.

Why this matters: Meta-layer education must update with the system.

3.12 Accessibility and modality gaps

Learning materials are available only in formats that exclude some participants.

Example: Onboarding depends on dense text and videos without captions, transcripts, screen-reader support, or low-bandwidth alternatives.

Why this matters: Education must be multi-modal and accessible by design.

Core Principle

Education in the meta-layer must be contextual, adaptive, participatory, and portable.

Participants should be able to learn across the web, grow with the network, and receive recognition for meaningful learning and contribution without surrendering agency to platforms, institutions, or AI tutors.

DP10 treats learning as a lived interaction between people, tools, communities, AI systems, and shared knowledge environments.

Example: A new participant enters a civic annotation zone. The interface explains what a zone is, offers a short guided task, introduces relevant glossary terms, shows examples from trusted community members, provides an AI learning assistant for questions, and awards a PEARL badge only after the participant prepares, engages, reflects, and applies what they learned.

What this feels like: You are not dropped into a system. You are accompanied into capability.

Without this: The meta-layer becomes another environment where insiders govern complexity and everyone else follows instructions.

Primary Mechanisms and Structural Conditions

5.0 Education Layer: Onboarding, Literacy, Guidance, Recognition, and Renewal

DP10 requires an education layer that makes the meta-layer learnable over time.

This layer includes:

- onboarding pathways
- interactive tutorials
- FAQs and guides
- shared glossaries
- AI learning assistants

- community knowledge maps
- peer teaching tools
- curriculum modules
- reflective learning workflows
- PEARL digital badges
- recognition of prior learning
- feedback loops for improving educational materials

The education layer must be embedded into the participant experience rather than isolated in separate documentation.

Failure mode: learning as an afterthought.

5.1 Onboarding and participant guidance

Onboarding must help participants understand what the meta-layer is, what they can do, what risks exist, and how to participate responsibly.

Onboarding SHOULD include:

- plain-language introduction
- guided first actions
- explanation of overlays, zones, bridges, agents, and credentials
- safety and privacy basics
- governance participation paths
- role-specific pathways (participant, educator, developer, steward, parent, student)
- progressive disclosure of complexity

Example: A participant first learns how to read a trust signal, then how to add an annotation, then how to join a governance zone, rather than receiving everything at once.

Failure mode: orientation overload.

5.2 Effective usage training

Participants need training that helps them use tools well, not merely access them.

Training SHOULD include:

- interactive tutorials
- sandbox environments
- scenario-based learning
- tool-specific guides

- “learn by doing” flows
- practice tasks
- error recovery guidance
- examples of good and poor use

Example: A bridge-building tutorial asks participants to compare sources, identify claims, attach context, and reflect on whether their bridge helps others understand the page.

Failure mode: feature exposure without skill development.

5.3 AI Learning Assistant

An AI Learning Assistant may provide personalized support, adaptive guidance, and real-time mentoring.

It SHOULD support:

- personalized learning plans
- context-aware explanations
- glossary assistance
- learning path recommendations
- comprehension checks
- reflective prompts
- accessibility adaptation
- multilingual support
- escalation to human educators or stewards

The AI Learning Assistant must comply with DP11 and operate within a visible capability envelope.

This includes:

- clear disclosure of AI identity, role, and limitations
- visible capability boundaries (what it can and cannot do)
- source attribution and provenance for educational claims
- expression of uncertainty and confidence
- escalation pathways for high-risk or sensitive topics
- restrictions in high-stakes domains (e.g., legal, medical, civic decisions)

The assistant SHOULD also be designed to avoid dependency by:

- prompting users to reason rather than only providing answers
- encouraging independent verification

- surfacing multiple perspectives where appropriate
- gradually reducing scaffolding as capability increases

Example: A learner struggling with “data sovereignty” receives a short explanation, a visual analogy, a community example, a prompt to apply the idea, and a suggestion to inspect a real data permission setting.

Failure mode: answer machine instead of learning scaffold or dependency engine.

5.5 Shared understanding glossary

A shared glossary is necessary to reduce misunderstanding across technical and non-technical communities. In the meta-layer, the glossary is not just reference material. It is **coordination infrastructure**.

A glossary SHOULD include:

- plain-language definitions
- technical definitions where needed
- examples
- related terms
- translations
- community-specific variations
- version history
- governance notes where meaning is contested

Terms likely requiring shared treatment include:

- meta-layer
- Metaweb
- Overweb
- overlay
- bridge
- zone
- smart tag
- presence
- agent
- credential
- reputation
- containment
- governance receipt

5.5.1 Term Object Model (Minimal)

Each glossary entry SHOULD be represented as a structured, versioned object:

- **term_id**: unique identifier
- **term**: canonical label
- **plain_definition**: accessible definition
- **technical_definition**: optional precise definition
- **examples**: usage examples
- **related_terms**: links to other entries
- **aliases**: synonyms or community variants
- **translations**: localized labels/definitions
- **zone_overrides**: context-specific definitions (see 5.5.3)
- **version**: semantic version (e.g., 1.2.0)
- **status**: draft, active, deprecated
- **provenance**: authors, sources, discussion threads
- **dispute_status**: undisputed, contested, under review
- **last_updated**: timestamp

Why this matters: Terms become interoperable artifacts that can be referenced, audited, and updated across systems.

5.5.2 Versioning and Change Management

Glossary terms MUST support transparent versioning.

- Changes SHOULD produce a new version with a diff
- Previous versions MUST remain accessible
- Breaking changes SHOULD be clearly marked
- Deprecations SHOULD include migration guidance

Example: “overlay” v1.1 clarifies scope; v2.0 changes definition boundaries. Interfaces show the active version and allow viewing prior versions.

Failure mode: silent semantic drift.

5.5.3 Zone-Scoped Definitions and Overrides

Communities (zones) MAY define contextual variations while maintaining a shared baseline.

- **global definition** provides baseline meaning
- **zone override** refines or constrains meaning in context
- overrides MUST reference the global term and declare differences

Example: “agent” in a medical zone may require stricter capability constraints than in a casual chat zone.

Why this matters: Enables local relevance without breaking global interoperability.

Failure mode: incompatible local meanings that cannot interoperate.

5.5.4 Governance and Authority Model

Glossary evolution requires governance.

Communities SHOULD define:

- who can propose new terms or edits
- review and approval workflows
- quorum or rough consensus thresholds
- escalation paths for contested definitions
- stewardship roles (editors, maintainers)

Decisions SHOULD be transparent and archived.

Failure mode: either centralized doctrine or uncontrolled fragmentation.

5.5.5 Dispute Resolution and Contestation

Terms can be contested.

Systems SHOULD support:

- attaching arguments and evidence to terms
- marking dispute status visibly
- parallel definitions when consensus is not reached
- time-bounded review cycles

Example: Competing definitions of “reputation” coexist with clear labeling and provenance until resolved or stabilized.

Failure mode: hidden disagreement leading to coordination breakdown.

5.5.6 Cross-Zone Translation and Mapping

Different communities may use different terms for similar concepts.

Glossary systems SHOULD support:

- mapping equivalent or near-equivalent terms
- indicating degree of equivalence (exact, partial, contextual)
- translation across languages and domains

Example: “annotation” in one community maps to “note layer” in another with partial equivalence.

Failure mode: siloed vocabularies that block collaboration.

5.5.7 Interface Integration

Glossary access must be embedded at the interface level.

Participants SHOULD be able to:

- hover or tap to see definitions
- view term provenance and version
- see dispute status
- switch between plain and technical definitions
- access examples and related terms

What this feels like: Understanding is available exactly where confusion arises.

Failure mode: glossary exists but is not used.

5.5.8 Feedback and Adaptation (DP18 Link)

Glossary entries SHOULD accept structured feedback (see 5.14).

This includes:

- clarity issues
- conflicting usage reports
- accessibility concerns
- translation gaps

Feedback SHOULD feed into version updates and governance review.

Failure mode: glossary stagnation despite participant confusion.

5.5.9 Safety and Misuse Considerations

Some terms may be targets for manipulation (e.g., redefining safety, trust, or authority terms).

Systems SHOULD include:

- audit trails for changes
- alerts for high-impact term edits
- review requirements for sensitive terms

Failure mode: semantic attacks that reshape governance through language.

Failure mode: language drift that fragments coordination.

5.7 Recognition of prior and informal learning

The meta-layer should recognize learning that occurs outside formal institutions.

This includes:

- peer teaching
- community stewardship
- annotation work
- civic mapping
- facilitation
- moderation
- translation
- technical contribution
- family and youth learning
- lived-experience expertise

Recognition systems SHOULD allow institutions to evaluate badges and learning artifacts without controlling the entire learning process.

Failure mode: informal learning remains invisible.

5.8 Cross-system credential recognition

Educational credentials must be portable across systems, communities, and institutions.

Credential frameworks SHOULD support:

- open schemas
- evidence bundles

- issuer identity
- verification
- translation to institutional credit or portfolio use
- expiration or renewal where appropriate
- dispute and correction pathways

Example: A PEARL badge earned in a municipal civic mapping project can be reviewed by a school, employer, or community organization as evidence of collaboration, research, and digital literacy.

Failure mode: credentials trapped inside one platform or community.

5.9 Formal curriculum integration

The meta-layer can support formal education without becoming dependent on it.

Curricular uses MAY include:

- sustainability projects
- citizenship education
- intercultural dialogue
- media literacy
- critical AI literacy
- research methods
- collaborative annotation
- public presentation
- SDG-aligned inquiry

Teachers and students should be able to use meta-layer tools for inquiry, annotation, co-creation, and public knowledge contribution.

Failure mode: the meta-layer is treated as an add-on tool rather than a learning environment.

5.10 Family-centered and intergenerational learning

Education should include home, peer, and community contexts.

Family-centered pathways MAY include:

- parent guides
- youth safety explainers
- family digital literacy activities
- educator referral materials

- community flyers
- library workshops
- intergenerational dialogue prompts

Example: A parent helps a child compare AI-generated and human-authored content using a simple overlay guide, then shares feedback with a school or community group.

Failure mode: learning remains locked inside institutions or expert spaces.

5.11 Community-authored knowledge environments

Communities should be able to create and maintain their own learning spaces.

These may include:

- knowledge maps
- annotated reading paths
- trusted source collections
- bridge libraries
- local civic guides
- peer learning circles
- community curricula
- project-based learning paths

Community-authored learning helps participants see knowledge as something they can contribute to, not only consume.

Failure mode: education becomes centralized content delivery.

5.12 Critical digital and AI literacy

DP10 must help participants understand synthetic media, algorithmic influence, provenance, and AI-generated content.

Learning SHOULD include:

- identifying AI-generated content
- evaluating source credibility
- understanding provenance
- recognizing engagement manipulation
- reading transparency cards
- interpreting trust signals
- understanding agent capabilities and limits

- knowing when to escalate or seek human expertise

Failure mode: participants use AI-rich environments without understanding synthetic influence.

5.13 Multi-modal learning access

Education must work across modalities and devices.

Learning materials SHOULD support:

- visual guides
- audio explanations
- transcripts
- captions
- screen readers
- mobile-first access
- low-bandwidth versions
- spatial or immersive learning where appropriate
- cognitive accessibility modes

Failure mode: educational access depends on one dominant interface.

5.14 Learning feedback loops

Educational systems must learn from participants.

DP10 should integrate with DP18 through structured learner feedback objects.

This includes:

- learner feedback objects
 - comprehension checks
 - tutorial effectiveness reports
 - badge appeal pathways
 - community review of educational materials
 - adaptation receipts when learning paths are updated
-

5.14.1 Learner Feedback Object (DP18-aligned)

A learner feedback object SHOULD include:

- **feedback_id**: unique identifier
- **learner_id**: participant reference

- **object_ref**: tutorial, badge, glossary entry, or learning artifact
- **feedback_type**: comprehension issue, clarity issue, accessibility issue, error report, suggestion
- **severity**: low, medium, high
- **description**: participant-provided input
- **context_snapshot**: state of learning interaction
- **resolution_status**: open, in review, resolved
- **response_action**: update, clarification, escalation, no action
- **reputation_weight**: influence based on participant trust level
- **timestamp**: submission time

Why this matters: Learning systems must evolve based on real participant experience, not only designer intent.

Failure mode: outdated or ineffective learning materials persist without correction or accountability.

Governance, Accountability, and Agency Surfaces

Education systems shape what participants believe the meta-layer is. They therefore require governance.

Communities SHOULD define:

- who can author official learning materials
- how materials are reviewed
- how glossary terms are updated
- how badges are issued and revoked
- how AI learning assistants are constrained
- how youth-facing materials are approved
- how community curricula are recognized
- how contested educational claims are handled

Governance should preserve both coherence and plurality. The meta-layer needs shared language, but it must also allow communities to teach from their own context.

Failure mode: education becomes either centralized doctrine or fragmented confusion.

Participant and Community Agency Surfaces

Governance of education is also a matter of agency for those being educated. Learners are not only subjects of curriculum decisions; they are participants whose comprehension, objections, and lived context are evidence about whether the material works.

Learners must be able to:

- see who authored a learning material, when it was last reviewed, and whether it is current
- see the version and provenance of any glossary term, including dispute status
- submit structured feedback on clarity, accessibility, or error, and see its resolution status (5.14)
- contest a badge decision, request review, and have the outcome recorded
- understand what an AI Learning Assistant can and cannot do, and reach a human instead
- learn in their own language and modality without receiving a degraded version of the material (DP21, DP23)
- decline data collection about their learning without losing access to the learning itself (DP4)

Communities and educators must be able to:

- author, adapt, and translate materials for their own context without seeking central approval
- set stricter conditions for youth-facing and high-stakes materials
- issue, review, and revoke credentials under published criteria
- audit badge issuance for farming, bias, and inconsistency
- retire or mark outdated materials and publish what changed

Institutions must be able to evaluate learning evidence without acquiring control of the learning process. Recognition should flow toward evidence, not toward whoever administers the platform (5.7, 5.8).

Failure mode: curriculum capture, where a single authority determines what counts as understanding and learners have no route to contest it.

Failure mode: feedback theater, where learners can report confusion but materials never change.

Community Signals Informing DP10

Signals from educators, parents, youth participants, community facilitators, and self-taught builders converge on a consistent picture: access to information has increased while support for developing understanding has not.

Recurring signals include:

- newcomers reporting that documentation assumes the knowledge it is supposed to provide
- participants who can operate tools but cannot explain the rights, risks, or governance around them
- communities discovering mid-conversation that they have been using the same words differently
- educators asking for materials that work in a classroom without a technical prerequisite
- parents and caregivers looking for accessible guidance on synthetic media and AI companions, and finding whitepapers
- learners with disabilities encountering onboarding that assumes dense text, video, and high bandwidth
- concern that AI assistants produce fluent answers while leaving learners less capable than before
- skepticism about badges, driven by experience with credentials that signal attendance rather than capability
- informal educators, translators, and peer mentors doing substantial teaching work that no system recognizes
- youth participants reporting that they are addressed as an adoption target rather than as future stewards

These are not requests for more content. They describe a missing layer: contextual, accessible, community-grounded learning with recognition that travels.

DP10 treats these signals as requirements. A system that cannot be learned by the people it governs is not public infrastructure, however open its specifications.

Evaluation Criteria

A DP10-aligned implementation should be evaluated against the following questions.

7.1 Learnability

- Can a new participant understand the system without expert help?
- Are learning paths staged and role-specific?

- Is complexity disclosed progressively?

7.2 Accessibility

- Are materials available across abilities, languages, devices, and bandwidth conditions?
- Are youth, families, elders, and non-technical participants supported?

7.3 Agency

- Does learning increase participant capability?
- Does the AI Learning Assistant teach judgment rather than create dependence?
- Do participants learn how to govern and contest, not only how to use tools?

7.4 Credential integrity

- Are badges evidence-backed?
- Are credentials portable?
- Can they be verified, corrected, or appealed?

7.5 Community grounding

- Can communities author and adapt educational materials?
- Are local examples and lived experience included?
- Do learning systems recognize informal contribution?

7.6 Currency and adaptation

- Are materials updated when tools, risks, or policies change?
- Are feedback loops active?
- Are outdated materials clearly marked?

Implementation Patterns

These implementation patterns translate DP10 into practical design moves, showing how onboarding, guidance, AI support, and credentialing can be embedded directly into participant experience rather than treated as external documentation or training.

8.1 Progressive onboarding paths

Introduce concepts in stages: presence, overlays, trust signals, annotations, zones, governance, credentials, and AI agents.

8.2 Scenario-based tutorials

Teach through real use cases such as evaluating a claim, joining a zone, submitting feedback, or earning a badge.

8.3 AI learning companion with guardrails

Use AI to adapt explanations and support reflection while preserving disclosure, sources, and escalation.

8.4 Living glossary

Maintain a versioned glossary with plain-language entries, technical notes, examples, and translations.

8.5 PEARL badge pathways

Design badges around Prepare, Engage, Reflect, and Leverage, with evidence required at each stage.

8.6 Community learning maps

Use knowledge maps to show relationships between concepts, artifacts, contributors, and learning paths.

8.7 Family and educator kits

Provide flyers, lesson prompts, safety guides, and simple explainers for home and school use.

8.8 Recognition bundles

Package badges with evidence, reflection, issuer identity, and verification for institutional review.

8.9 Learning retrospectives

Publish what participants struggled with, what materials changed, and why.

8.10 Accessibility-first publishing

Every core educational artifact should have accessible alternatives.

Lifelong Learning Opportunities

The meta-layer should support continuous learning across stages of life and participation.

Learning MAY include:

- digital literacy
- AI literacy

- governance literacy
- media and provenance literacy
- data sovereignty
- civic participation
- annotation and bridge-building
- knowledge mapping
- community facilitation
- developer education
- regenerative systems thinking
- professional skill development

Learning should be available inside meta-communities, not limited to formal courses.

Failure mode: education ends after onboarding.

PEARL Digital Badges

PEARL badges recognize learning as a process, not a single completion event.

PEARL stands for:

- **Prepare:** orient, understand context, identify goals
- **Engage:** participate in a meaningful task or community activity
- **Reflect:** demonstrate learning, insight, or changed understanding
- **Leverage:** apply learning to a contribution, portfolio, project, or next step

PEARL badges SHOULD be:

- evidence-backed
 - portable
 - revocable or correctable where necessary
 - linked to learning artifacts
 - interpretable by communities and institutions
 - resistant to badge farming
 - aligned with DP18 reputation and feedback systems
-

Minimal PEARL Badge Schema

A PEARL badge SHOULD include:

- **badge_id:** unique identifier

- **learner_id**: participant reference
- **issuer_id**: issuing community or entity
- **domain**: skill or knowledge area
- **level**: introductory, intermediate, advanced
- **pearl_stages_completed**: Prepare, Engage, Reflect, Leverage evidence
- **evidence_bundle**: links to artifacts, annotations, reflections, or projects
- **assessment_method**: peer review, instructor review, automated checks, hybrid
- **feedback_refs**: linked DP18 feedback objects
- **reputation_impact**: contribution to participant reputation
- **timestamp**: issuance date
- **expiration_or_review**: renewal conditions if applicable
- **verification_signature**: authenticity marker

Example: A participant earns a PEARL badge for “Provenance Literacy” after completing a guided inquiry, annotating a contested claim, reflecting on source quality, and applying the skill in a community context.

Failure mode: badges as decorative rewards rather than evidence of learning.

Relationship to Other Desirable Properties

DP2 – Participant Agency and Empowerment

DP10 strengthens agency by making the meta-layer understandable and usable.

DP3 – Adaptive Governance

Participants must learn how governance works in order to participate meaningfully.

DP7 – Interoperability

Learning artifacts and credentials must move across systems.

DP8 – Community-Defined Participation Zones

Zones need their own educational materials and onboarding paths.

DP9 – Developer and Community Incentives

Learning contributions, teaching, translation, and curriculum work should be recognized and incentivized.

DP11 – Safe and Ethical AI

AI learning assistants must be disclosed, bounded, accountable, and contestable.

DP14 – Trust and Transparency

Education makes trust signals interpretable.

DP18 – Feedback Loops and Reputation

Learning feedback, badge integrity, and educational reputation depend on DP18.

DP19 – Amplifying Presence and Community Engagement

Education converts awareness into durable participation.

DP21 – Multi-Modal Interactions and Experiences

Learning must be available across modalities and assistive contexts.

Non-Goals and Explicit Boundaries

DP10 does not:

- prescribe a single curriculum, pedagogy, learning platform, or credentialing authority
- require formal education, enrollment, or institutional affiliation to participate in the meta-layer
- treat credentials as a prerequisite for rights, access, or standing
- guarantee that learning produces agreement, or treat disagreement as a comprehension failure
- replace schools, libraries, universities, employers, or professional certification
- mandate AI-assisted learning, or prohibit it
- require participants to disclose learning history, difficulty, or progress as a condition of access (DP4)
- assume literacy, bandwidth, language, ability, or device access as a baseline
- promise that all learning can be evidenced, or that unevidenced learning is invalid

DP10 defines the conditions under which learning remains accessible, adaptive, community-grounded, and portable. It does not define what participants must know.

Minimum DP10 Alignment (Non-Normative)

Minimum alignment is not a feature checklist. It is the threshold at which a system can be learned, questioned, and grown into by the people it affects, rather than only by those who already understand it.

A DP10-aligned implementation should, at minimum:

- provide a plain-language entry path that explains what the system is, what participants can do, and what risks exist, before requiring any consequential action (5.1)
- disclose complexity progressively, with role-specific pathways rather than a single undifferentiated onboarding flow
- publish a versioned glossary with plain-language definitions, provenance, dispute status, and access at the point of confusion (5.5)
- disclose AI learning assistance clearly, bound its scope, cite sources for educational claims, express uncertainty, and provide escalation to a human (5.3, DP11)
- back any credential with evidence, and support verification, correction, revocation, and appeal (**PEARL Digital Badges**, 5.8)
- make credentials and learning artifacts portable across systems and legible to institutions without transferring control of the learning process (5.7, 5.8, DP7)
- provide accessible alternatives for every core educational artifact across ability, language, modality, device, and bandwidth (5.13, DP21, DP23)
- accept structured learner feedback and connect it to material revision with visible resolution status (5.14, DP18)
- teach governance and agency alongside tool use, including how to contest decisions and participate in rule-making (DP2, DP3)
- mark or retire outdated materials rather than leaving stale guidance in circulation
- avoid conditioning access to learning on data collection about the learner (DP4)

Partial compliance that provides tutorials without accessibility, badges without evidence, AI assistance without disclosure, or feedback without revision should not be treated as alignment. Each of those omissions reproduces the failure DP10 exists to prevent: a system that can be operated but not understood, and therefore not governed.

Open Questions and Future Work

1. What minimum glossary terms should be standardized across the meta-layer?
2. What schema should define PEARL digital badges?
3. How should informal learning be verified without institutional capture?
4. What evidence should be required for different badge types?

5. How should AI Learning Assistants disclose limits and sources?
 6. What youth-safety standards should apply to educational materials?
 7. How should community-authored curricula be reviewed and versioned?
 8. What credential translation tools are needed for schools and employers?
 9. How should learning feedback objects integrate with DP18?
 10. What standards ensure accessibility across modalities?
 11. How should outdated educational materials be marked or retired?
 12. What forms of learning should affect reputation or role access?
-

Path Toward ML-RFC

DP10 is currently an ML-Draft and serves as exploratory scaffolding for education, onboarding, credentialing, and lifelong learning in the meta-layer.

Advancement toward ML-RFC status SHOULD require:

- a minimal shared glossary standard
- a PEARL badge schema
- evidence bundle requirements for credentials
- AI Learning Assistant disclosure and safety requirements
- accessibility requirements for educational materials
- learning feedback object standards
- curriculum versioning and review processes
- cross-system credential recognition pilots
- youth and family learning safeguards

Early ML-RFC candidates may focus on:

- Shared Meta-Layer Glossary Standard
- PEARL Badge and Evidence Schema
- AI Learning Assistant Safety Requirements
- Learning Feedback Object Standard
- Credential Recognition and Translation Framework

DP10 will likely mature through multiple component RFCs rather than one monolithic standard.

Closing Orientation

DP10 makes the meta-layer learnable.

It ensures that participants are not merely onboarded into tools, but accompanied into agency, literacy, stewardship, and contribution.

A DP10-aligned meta-layer teaches people how to see the web differently, how to act with confidence, how to recognize trustworthy context, and how to grow with their communities.

Onboarding becomes orientation.

Education becomes participation.

Badges become evidence of growth.

Learning becomes shared infrastructure.

This is how the meta-layer becomes not only usable, but teachable, transmissible, and alive across generations.
