

Intermedia Markup Language

an [Inclusive Organizing](#) module



This work is licensed for sharing under the [Creative Commons Attribution 4.0 International License](#).

This is a [living document](#).

For a static version, use the [File](#) options to [Make a Copy](#) or [Print a PDF](#), or [email](#) to discuss potential adaptations.

Note: Intermedia Markup Language is already licensed for sharing, but several sections are still incomplete as of March 2025. This document is high priority-- but not top priority, within IO's dozens of modules-- for further updating.

Preface

Purpose

Intermedia Markup Language (InML) is a markup language for developing socially important, interoperably linked data and metadata (including resource description frameworks) regarding verbal expressions and media items, as well as structurally-defined selections within media items. This type of markup language can facilitate the development, stewardship and curation of richly informative media networks, directories, repositories and wikis, in both governed and ungoverned networking contexts, thereby greatly enhancing the interoperability and navigability of cross-organization and cross-network data. It can provide root grammars, or pattern languages, for developing highly distributive and modularized digital networking, including technical integrations between different networking systems and tech stacks.

Formal Context

InML is a primary component of [Inclusive Organizing](#) (IO), and can function as a subsystem for other IO components. InML is especially meant to function as a subsystem for fully developing a [Intermedia Descriptive Framework](#).

While intermedia markup language can essentially be an open protocol, and can act likewise as a supplement to [Foundations](#), a key premise in InML is that media markup is only fully and fairly, accurately and consistently scalable in governed contexts.

Note that InML is meant to be used in the context of specific communities, although some of those communities could (if sustainably supported) be defined quite broadly; for example, global communities of millions to billions.

Overview

InML can enable all authorized media users to annotate specific media items, and selections thereof, for these purposes:

- **Assigning metadata tags:** *Personally assigning (binary or dimensional) [attributive tags](#), including semantic links*
- **Evaluating tag assignments:** [Evaluating](#) tags which have already been assigned
- **Structured conversation:** Attaching (potentially threaded or [structured](#)) verbal commentary/conversation

Note that responsible InML usage will often depend upon project and network governance systems, including their media permissions or capabilities systems which provide viewing, editing and/or annotation authority to users within identified cultural contexts. InML can also be used as openly as (for example) the Internet Protocols, but most of its potential is based on social accountability. For example, InML annotation authority can be associated with governed participation roles such as the *primary roles* in [Inclusive Governance Framework](#). Such governed roles could enable securely-identified agents to responsibly develop socially important media metadata across many networks representing diverse communities.

Media Sharing Types, Sources and Governance

Digital media items are shared in several basic ways:

- distributing (copying and transmitting)
- quoting/addition (partial or complete copying)
- linking (*hypertextual reference*)
- transclusion (*hypertextual incorporation*)
- paraphrasing (which is inherently editorial and creative)

Each type above creates a different data infrastructure and thus has-- or, at least, *should have*-- a fundamental effect on how we describe and tag media items. *Add some explanatory text*

Each media-sharing type (and, especially, complex combinations of them) can have relationships with multiple media sources and governance structures, and can be governed as-needed according to official authorizations for signal sending, receiving, accessing and editing. (See [Intermedia Resource Description Framework](#) if desired for context regarding authorizations.)

Types of Expression

This section provides several suggested “primitive” types of expression. which can be applied to all utterances and written language.

Primitive Expressive Types

The following primitive or root expressive types can be applied to all utterances and written language. All expressions can be usefully defined as containing at least one relationship, concept or [semantic triple](#), although one or two elements of that triple may be unstated but contextually *implied*. These primitive expressive types are defined as relationships.

Indicative/ descriptive

Indicative statements express percepts or concepts, indicating one or more relationships or probabilities which are perceived (or simulated) in some supposed or theoretical context.

- Indicative statements may refer to perceived types or instances of elements, forms and actions.

- Indicative statements may distinguish *correlation* from claims of *causation*
- SUBTYPE [realis](#)/ declarative
- SUBTYPE, INDIRECT: Figurative
 - sub-subtype: Analogous/ metaphorical: (partially) describes a relationship by indicating a supposedly similar or identical *type of relationship*
- SUBTYPE: Simulation
- SUBTYPE Fictional/ whimsical
- *Potentially allow anyone to apply tags, but distinguish self-tags*
- Suggested collective feedback for indicative statements: see [Evaluative Scales and Reaction-Types](#)

Inferential

- Inferential statements include **deductive or inductive** speculations about relationships or relationship-types which have not been directly perceived.
- These statements may (or may not) indicate perceptions of confidence, uncertainty and probability.
- SUBTYPE: Propositional
 - Describes something which *seems* possible, and may be explored, without expressing a position regarding its probability.

Note: Semantically directional indicative and inferential links address the fundamental distinction between **unique** and **generalizing/typifying** relationships.

- Recipe, Plan, Report and Model/Map should be primary recommended tags for indicative statements
 - For example, participants should be able to tag traits such as "declarative observation/report" or "declarative abstraction"

Interrogative/ questioning

- Interrogative statements (including all questions) *request feedback* from others

Instructional/imperative

- Note: May be coercive “orders”, or mild suggestions, or anything in between. (We also may make indicative and interrogative statements in explicitly *coercive* or *threatening* contexts, but that doesn't make them imperative statements in their fundamental structure. They could however be tagged as implicitly coercive.)

Functional Expressive Types

Some very broadly relevant media tags are suggested below. However, *each collective should be free to collaboratively generate its own preferred tags*. Also, tag usage should be optional and may often seem unnecessary and inefficient for relatively simple, brief discussions.

definition

protocol

essay

discussion

comment/ commentary

Comments are directly linked to media items; *commentaries* may or may not be.

reply (linked to comment)

tangent

request

request for comments or editorial feedback

request for work (or participation)/ Want

counter-proposal

especially optional; for relatively detailed official negotiations

offer of work

question

issue/concern

Governance Resource Types

see [Foundations of Modular Organizing/Functional Media Primitives/governance resource](#)

Descriptive Tags

One of the main goals of intermedia markup language is to enable more informative, legible and *navigable* “social networking” feedback on the media items within any specific network.

Each collective should determine which media feedback and rating tags are available for use in their media directory. The options (which each collective enables or disables) should include all potential options for rating or replying to *comments* which have been attached to media items, or have otherwise become part of a [structured conversation](#) system. Collectives should also be able to give media contributors the *option* to toggle any of their collectively-supported media feedback options off or on for media they submit.

Note on Dimensionality

All attributive tags, evaluative scales and reaction types below are implicitly based on perceptions of dimensionality. Dimensional description is of these basic types: *nominal and scalar (ordinal, interval and ratio)*. Attributive tags and reaction-types are always *nominal* although they can include other (scalar) data; evaluative scales are scalar.

See *Traits* in [Foundations of Modular Organizing](#) if needed for further consideration of this topic.

Attributive Tags

It should be possible (and *optional*) for the authorized users in any specific network to tag media items and selections *of* items, including discussion topics, by indicating their personal perception of one *or more* attributes present in that media. Media attributes (tags) could be simple or structurally complex. Each tag should provide perceived information about the media item or selection thereof which they're attached to.

Applied attributes could be used to indicate that media items are observations or (more seriously and officially) reports. Reports could potentially be submitted anonymously, but hold no social accountability unless they are attributed directly to digital accounts.

Viewers should be enabled to easily *endorse or dispute* reports. --

- Also advocate the ability to designate an observation explicitly as a **report**, and to designate especially official and important reports as **testimony**.

Attributes are relative to the people who apply them; thus, any endorsement of assigned attributes should be associated with digital identities.

Many tags may be simple terms (or *dimension:value* evaluative pairs) which acquire all of their local meaning/context because of their precise placement within a specific media directory. However, verbally defined attributes can be made explicitly, structurally *relational*, by including context-creating semantic structure/grammar within them. For instance, an attribute may state "this item has Attribute X in relation to Link A" or "this item has attribute X in relation to Collective B". One key type of potential semantic attribute is the *semantic link*, which creates a contextualized link to another media item or selection. I.e. the link should allow attributes to be applied to the relationship

between the media items. -- some very useful link attributes: *similarity, conflict, relevance, more detailed or more general* .

Relational and Modular Tagging of Sections

To achieve its full communicative potential, InML should distinctly allow for *sections of media items* to be tagged either ‘independently’ or explicitly according to their perceived functional value in relation to (1) the item they’re part of or (2) another specific item. -- Or, in the case of modular systems (potentially including ‘Modular Requests’), selections of items may be rated according to their perceived value as a specific component in a modular system.

Evaluative Scales and Reaction-Types

Many evaluative scales and reaction-types should be of interest to most projects and communities, which can also generate new scales to reflect their unique priorities. This document emphasizes scales and reaction-types of general interest.

Well-developed scales and reactions will come much closer to representing the boundless (and dynamically evolving) complexity of our social and cultural attitudes. Most efforts to approach such complexity have been through in-depth polling based on evaluative scales, in stark contrast to the crude abstractions which are prevalent in social media reaction systems.

WIP: Emphasize five point scales which enable agents to report ambivalence (see [Covalence](#)) with declarative statements and propositions-- including commentary--and allowing indications of perceived relevance, importance, priority and confidence/uncertainty.

Briefly describe and link, the use of [Covalence](#) (and Collective Covalence, in that link).

Also note the potential of [Consensus Calculus](#) to reduce insincere [strategic voting](#) or ratings in large groups, and especially in anonymous polls. (This author recommends consent-based decision policy for all non-emergency group decisions; however, ratings systems can & should be used as part of discussion and design processes, and it’s important to make them as “un-gameable” as possible.)

Suggested peer feedback metrics are listed below:

Agree/Disagree

3 to 10 point scales, centered on apathy or ambivalence. As a five point scale: Strongly Agree* Agree* Abstain* Disagree* Strongly Disagree

- *Agree/disagree* is only relevant, technically, for [indicative](#) (often declarative) statements. However, communities should be able to sort that logic out among themselves. (Some communities may wish to analyze and explicitly tag “Declarations”.)
- People should be able to change their positions at any time, but each person’s history of positions (and the history of media editing) should be easily accessible.

Endorse/Dispute

3 to 10 point scales. As a five point scale: Strongly Endorse* Endorse* Abstain* Dispute* Strongly Dispute anything tagged as Report/ Observation.

- People shouldn’t directly dispute reports unless they have directly contradictory experience and/or evidence.
- People should be able to change their positions at any time, but each person’s history of positions (and the history of media editing) should be easily accessible.
- Among other things, this rating should link directly to [Skill Standards](#) (for rating skills, whether they have or haven’t been self-assigned or otherwise attributed yet)

Relevance

3 to 10 point scales. As a five point scale of (personally perceived) Relevance (to the topic or thread in question): Highly Relevant* Relevant* No opinion* Irrelevant* *Highly irrelevant

- *WIP: integrate the “Relevance” concept with a “Tangent” and/or “Fork” grammar???*

Importance

3-10 point scales.

- Importance (along with [priority](#)) is a very useful concept within reasonably well-defined contexts, such as project management and structured conversation. See [Developmental Staging Recipe](#) for further context.

Priority

3-10 point scales.

- Priority (along with importance) is a very useful concept, but only within reasonably well-defined contexts such as project management. See [Developmental Staging Recipe](#) for further context.

Trust

(placeholders below; not integrated yet)

Trust is the [dimension](#) of [confidence](#) in an [agent's probability](#) of [performing](#) a [desired action](#) or [activity](#) within a [specific context](#).

(I.e. trust occurs in relation to relationships between the trusted (or untrusted) agent with other agents, resources and/or events)

- *integrate trust-level conceptions such as*
<https://blog.discourse.org/2018/06/understanding-discourse-trust-levels/>

eventually (not necessarily within ‘trust’) also integrate the four player archetype model:
https://en.wikipedia.org/wiki/Bartle_taxonomy_of_player_types

Distrust

Distrust is the [dimension](#) of [confidence](#) in an [agent's probability](#) of failing to [perform](#) a [desired action](#) or [activity](#), or of performing a feared activity, within a [specific context](#).
(*I.e. trust occurs in relation to relationships between the trusted (or untrusted) agent with other agents, resources and/or events*)

Product Rating

Product ratings are for [goods](#) and/or [services](#).

- **3-11 point scales. As a 5 point scale of 1 through 5: Excellent* Above Average* Average* Below Average* Poor anything tagged as a *Product* (Good or Service).**
 - Note that 10 point scales, while convenient in decimal systems, only have a neutral rating as long as 0 is allowed. Such scales are recommended for their situational subtlety and their excellent compatibility with decimal-based mathematics.

Product ratings may relate closely to [trust](#) (of product-providing agents), but it's different. (Ratings of services are, essentially, a precise subtype of trust.)

Rubrics

Note that academic [rubrics](#) are interoperably valuable for other product rating contexts, given that teachers put vast amounts of professional attention on such assessment techniques. Some suggestions for modular rubric development:

1. Rate key [components](#) and (separately) key [traits](#).
2. For highly detailed rubrics, rate *traits of (some or all of) the key components*.
3. Assign an overall rating (score/grade) which is NOT mathematically calculated, but should make sense given the ratings of components and traits, *and*
4. **Track all partial ratings which are especially low or high, across any important social context**, such as (in academics) a course or program. I.e. perhaps Alfred gets good or excellent overall scores generally, but has a pattern

of getting "needs improvement" for the same trait or the same type of component across many documents.

Emoji

- Optionally, add diverse emoji "reactions" (like Slack & Discord, etc, "reactions") which will be visible across a clearly identifiable social context.

Flags

- **Flags** can be enabled to indicate concerns regarding collective policies and potentially restricted content.
 - See [Inclusive Governance Framework](#) and [Consent-Based Governance](#) for deeper consideration of collective standards and restrictions.

Currency Transfers

- **Collective currency distributions can be enabled per open patronage/funding requests, and “bounties” for fulfilling open media development requests**
 - This is distinct from collective mutual credit transfer agreements
 - Each collective or any (potentially temporary, makeshift) subgroup thereof, may make agreements to fund specific media Requests-- and in receiving responses to such Requests, may distribute according to a collective-specific value equation (derived from an ongoing generic global model/recipe) which awards proportionally to all well-ranked responses. Though, in cases where a single ‘solution’ must be eventually chosen (by an individual or a team), there should perhaps be a bonus given to the final choice.)
 - Distributions for responses/‘solutions’ to funded Requests should be made proportionally based on the ratings BY THE FUNDERS of the responses.

- Bear in mind that the Request may be simply for someone to add a type of media/ informational resource to the collective commons-- and be funded by a group of independent individuals-- OR be a Request to perform a creative task for a specific team-controlled Process. If it's for a team-controlled process, the funders should make their decisions of which response(s) to use ****as a team****, instead of proportionally per person.
 - Ideally, distributions will be as finely/ gradiently based upon each valued contribution as possible. For instance, if someone adds an open source idea which does roughly 90% of what the Request required, but then someone else adds the rest of the needed contribution, then the 'someone else' probably deserves the full last 10% of the funding. -- However, this is a new realm, which communities must explore together. Also note that as with all collective processes, the collective should be protected against "rush jobs". I.e.: a Request should normally indicate a (reasonable) time frame during which responses can be openly added by anyone, probably before any financial distributions are made.
 - Also bear in mind the concept of Modular Requests, and how that could affect distributions.
 - Also bear in mind that people may make post-facto distributions to valued contributions if they become wealthier in the future.... HOWEVER, discourage complicated mandatory value equations!
- In addition to that: **If** a collective has any collective-wide advertising revenue (from advertisers which have been approved by collective-wide agreement), much/ most of that advertising revenue may be broadly distributed according to any collective-specific formula involving positive feedback AND/OR a strong ratio of positive feedback per item to each item's total views.

- Users within a particular conversational context, or topic, should have options to sort and filter content according to whichever peer feedback metrics have been used there.

Important added note: It should also ideally be possible to make comments *and* to apply peer feedback mechanics to others' contributions either openly (within a given thread) **or privately**. (Open feedback is good practice generally IMO, but private interactions are valuable, and ought to have powerfully precise tools.)

Semantic Co-Rating

Add a brief overview of [Semantic Co-Rating](#) as the logical evolution of massively scalable ratings format options:

Items do not get rated.

Tags do not get rated.

User profiles don't get rated.

Things that get rated:

Tag(s) + item(s) (relevance)

Item(s) + tag(s) (rating)

Linked Commentary

Comments can easily be (and often are) linked to specific media items, either in unstructured chronological sequence or with subthreading. Linked comments enable expressivity and nuance which isn't strictly possible with attributive tags or evaluations.

The greater potentials of linked commentary can be developed with [Structured Conversation](#).

Structured Conversation

Preface

This is a prototype for an open source protocol to enable the generation of *complexly and compoundly structured* conversations with potentially collaborative inputs by any number of persons.

The prototype should enable inter-networked media hosts *and* users to link written commentary to digital media resources *and* to (technically identifiable, targetable) *selections of media items* which are stored anywhere.

This model should enable participants in any media network to attach commentary to any media resources which they technically can comment upon, either because commentary is ungoverned or because they've been authorized (either personally or via assigned roles/capabilities) by the media creator(s) or by authorized sharers of the media in question-- and, to *evaluate commentary* by explicit collective-based evaluative standards.

This model should enable people and communities to generate truly complex nonlinear dialogue related to all types of recorded, asynchronously shared digital messages and media. For instance, media contributors and users should be able to generate *recursive* written conversational structures, in which the same user interface should be (automatically) enabled at every (potentially hierarchical/*nested*) level of complex interactions.

Goal: This structured conversation model is meant to help people reduce the extreme inefficiency and incompleteness of most online written discussions of specific topics.

Relationship to Intermedia Markup Language

This Structured Conversation model is part of Intermedia Markup Language, creating a conversational function which supplements the [Media Description Tools](#). Media feedback tools and conversation structuring tools can be used separately or together. I suggest that their potentials are closely related, and while they can be trialed and implemented separately, they should ideally be used together.

- **Collective conversations may or may *not* be directly attached to established media items.** (In many or most cases, they will be attached to discussion channels or forums instead of media items, and will be identified and sorted according to [Subject](#).)

- A structured conversation model can apply to both live/streaming conversation (including in-voice and in-person) and to conversational exchanges of written messages.

Basic structured conversation format

Attachment

Structured conversations should be attached at the top level directly to media items, just as attributive tags are.

Subject

Each level of discussion (an intentional level of design) has a specific *Subject*, which is either a media item, a written topic attached to a media item, or a “standalone” written topic. A *thread* (a serial stream) of comments can be attached to the Subject.

Subjects may include any media items, or any *directly selected elements* of media items, in any technically compatible format. For example:

- **Text, links and/or inline files** posted intentionally to receive commentary (and content ratings) in a media forum.
- **Comments, or *selected parts of comments***, within discussion threads.
- Note: technically, *Comments* is a type of tag/attribute, which would link to a separate stream.
- should offer the same options for every topic (or subtopic) within a specific collective; per each media item.
- Must integrate a general “conversation for action” recipe
- Integrate “Meetings”, and note that the art of *facilitating meetings* is distinct from the art of *facilitating objectives and processes*
 - Meetings should often be oriented toward complex brainstorming and discussion, with deliberative statements, proposals and decisions left to unhurried and inclusive asynchronous media (such as Loomio, but with many more options.)
- The suggested elements of structured conversation should directly integrate with descriptive tag types in InML and related protocols.

Conversation Threading

The attachment of a comment to any media item (*including* an existing comment) without attached comments should create a new linear commentary level/thread. (The new comment will be the first entry at that level.)

Within each commentary level/thread, undirected comments will be added by default in a linear chronological sequence, with new entries added at the top. However, communities should be able to offer a wide variety of comment-sorting options for their users, based on the content-classifying and rating mechanics they use.

Within each thread, it should be possible (and *optional*) to:

- to **title** each comment
 - titles *may* serve as brief overviews or “TL;DR” summaries of long comments
- to **direct** each comment as a *reply*:
 - as a reply to the topic of that thread
 - as a reply to one *or more* of the specific comments in that thread
 - directed comments should display a clear visual link to whatever they’re replying to.
 - this visual link may be constantly visible, *or* may only show up when one mouses over the comment’s “in reply to ____” marker.
- to **tag** each directed comment, indicating specialized function/purpose:
 - question
 - answer (to someone else’s question)
 - issue/concern

Ideally, all comments in a specific conversational level/thread should be strongly related to the topic of the thread, *or* to one or more comments in that thread.

Significant tangents can be opened up as ‘subthreads’ (focusing on a clearly definable *element*/subtopic of the thread’s topic) or as new topics, either of which create new nodes in this system.

- It should be up to commentators whether or not they view their targeted replies to *other comments* as minor tangents/responses, or as new subthreads. (So, they should have an option to reply on the same level, or or to create a new level.)

However, communities will of course need to work together on a shared sense of best practices.

If individuals wish to rate or reply to *part* of an existing topic or comment, they should be able to highlight/select *any part(s) of the text*, and:

- comment directly upon that selection, within that thread
- apply *collectively-selected and user-selected* peer feedback ratings to that selection (see [Media feedback mechanics](#))
- create a new subthread, focused on that selection

Note: For this intentionally complex system to work, it must make it easy for people to reference and visually link to other comments in the same stream/thread. -- and, people should also ideally have visual indicators if they're currently drafting keywords phrases and/or entire passages identical or very similar to previous content in the stream. (this is advanced but very powerful potential)

Semantically Limited Subthreading

The boundlessly recursive "reply" feature in many forums and networks, including Reddit, is often inefficient and illegible. New subthreads should be created *only* according to genuine intentional tangents or 'sharpened focus' of the current topic. Such tangents or focuses are only occasionally a part of back-and-forth discussion. *However*, the question of whether to reply in a thread or to create a subthread will be mostly a matter of culture and collective standards, and not a technical issue, *if* discussion tools are made properly analytical and recursive.

- IF you are replying DIRECTLY to the WHOLE reply of someone else who is replying to the WHOLE media item, it should still be a subthread. HOWEVER you can generalize your reaction/ response to another's comment to make it a legitimately autonomous comment on the main thread-- and you may want to tag them/ credit them for that...
- There's a problem here because we're not usually disciplined in indicating when we are replying to PART of a media item instead of the WHOLE. Creating consistently rational organization of discussion (to the limited degree that that's possible) will require that distinction... however, it should not be 'forced' on

people. If people desire to riff creatively on the parent thread of a given media item, they should be allowed to brainstorm: but, if others desire to move forward discussion in an orderly analytical way, they should have the tools to do that, AND to receive (perhaps positive) feedback from peers for doing so.

In addition to having the ability to focus better on specific simple or complex ideas, questions and concerns, discussants in **this system should generate a MUCH better picture of precisely what people do and do not agree upon.**

With some **effective visual mapping** (crucial!!!) **it should also be relatively easy for viewers to quickly spot specific points of high agreement AND high conflict** within any complex discussion.

-- It must be made easy for people to reference and visually link to other comments in the same stream/ thread. -- and, people should ideally have visual indicators if they're currently drafting keywords phrases and/or entire passages identical or very similar to previous content in the stream-- in a clickable way. (this is advanced but powerful)

Meetings and Messages

Structured conversation can be organized according to a basic "meeting or message" distinction.

- See [Inclusive Meeting and Event Framework](#) if desired regarding meetings in general.)
- A **meeting model/recipe/protocol can be functionally related to a general model for [Goal-oriented Structure](#)**. (Noting that meetings aren't *only* useful for "conversation for action"; they may include unfocused coworking sessions as well as recreational activities.)

Collective members should be easily able to schedule live meetings (events) for specific groups or subgroups in any format for any reason, with strictly defined or open-ended agendas. Members should be enabled to schedule live in-person meetings; however, inclusive meeting opportunities should increasingly include telemeeting/ remote participation options per developing tech such as [WebRTC](#).

One of our main social shortcomings is, arguably, the current deficit of very short and well-defined live meetings to address minor but confusing issues. Defining “micromeeeting” practices effectively within a general networking context could alleviate that.

Goal-oriented Structure: Modular Organizing Recipe

Intermedia Markup Language and *Structured Conversation* are closely related to [Modular Organizing Recipe](#) (MOR) and its components, such as [3D Development Recipe](#) (3DD). 3DD could help people and teams to develop InML standards and media networks for specific communities.

3DD is somewhat analogous to Conversation for Action (CfA) models. CfA focuses on offers, requests, negotiations and commitments in low-trust networks and markets. 3DD includes such transactions, but is focused on team-based discussion, design and decision processes which may (or may not) use [Consent-Based Governance](#) or other consent-based decision standards. (For example, [Project Management Tags](#) could be directly used in collective-specific decision protocols.)

Each of these p2p networking models and protocols should be separately usable. Consent-Based Organization is already [openly licensed](#), and the rest of these models will also be openly licensed when they seem reasonably stable and reliable for general use.

Media and Metadata Editing and History

Media contributors should be able to edit/revise/update all types of media, per each collective’s standards; **however**, content ratings should ideally only be *directly attached* to the exact file/item versions which they were based on. (A version is a *media resource* which has a *specific accountable history of editorial governance*.)

- This creates a slightly thorny dilemma: should we be allowed to significantly revise text files which others have directly commented verbally upon, or [evaluated](#), without deleting those comments or evaluations? -- But, as with so many other things, this really should ultimately be up to communities.

- Media revision histories, and feedback and comments for prior versions, should be easily accessible-- and, may be indicated via:
 - a. “visual shorthand” attached to the most recent/active media item version,
 - b. Significant detail via mouseover,
 - c. Fully (of course) by clicking on a prior version and making it the key entry in the current visual representation. (Also, ideally, enabling side-by-side comparisons of two or more entries.)

Evaluation of Revisions

Ideally, it should be possible for a collective to directly rate the *addition* of a media item to a resource base according to collectively-defined metrics, and also to directly rate *revisions* of media items (or collectively-recognized *versions* of those media items).

“Revision-rating” could be especially useful for media items which already are part of a collective’s resource base, because a collective is likely to be relatively familiar with those items. However, if edits aren’t highly technical or conceptually complex, it’s theoretically possible for participants to effectively rate edits to items which they haven’t seen before, according to explicitly described collective standards.

InML Networking Standards

Intermedia RDF

See [Intermedia RDF](#) regarding general intermedia hosting, editioning and versioning.

InML Record Creation

Each authorized digital identity should be able to contribute media item links of any type to create new records in a InML-mapped directory and, when applicable, add entire media items (files) into a related repository. These linked or added media items may be original statements, such as intended discussion topics, as well as anything else we typically think of as a media item. New records and files can go by default into an unsorted “Inbox” for items which haven’t been marked up with metadata or linked data. InML can then be attached through formal use of resource description frameworks

and/or any interfaces which enable people to attach one or more attributes to each media item's record.

Signal Security, Access and Editing Authorizations

To support the sustainable generation of governed media directories, InML must enable the use of one or more *securely networked* digital identity sign-in systems.

Digital signals-- both live streaming and asynchronously transmitted items/ packets, of all sorts-- should by default document which identity each signal has come from, in order to support social accountability **and** fluid intercommunity creativity.

Customizable viewing and annotation authorizations: Users should be able, in their *self-hosted content* and in their *media sharing to receptive communities*, to allow open participation (descriptive tagging, evaluation and/or commentary) by anonymous guests *and/or* any number of pseudonymous or legal identity accounts. But they should *also* be free to *limit participation* according to any standards which each collective allows, and which each media-contributing agent directly enables.

Private messaging: The collective markup model should enable direct *private* p2p messaging to individuals and groups, (via accounts linked to specific media item records, including tags, ratings and comments) according to *personal private messaging permissions*, with private messages and requests **directly attached** (when desired) to (users' personal versions of) those items or targeted *selections within those items*. This direct messaging should directly or indirectly enable live / streaming channels of communications in any technically feasible formats, including video.

I recommend that networks/intentional collectives use a consent-based model such as [Inclusive Governance Framework](#) to establish participation roles, rights and responsibilities per each digital identity and-- when desired, and manageable- per each unique individual. (Enabling multiple identities per person can be good for conversation and for many interactions; not so good for evaluative ratings.)

About Notifications

- Communities should be able to offer flexible options for users to receive separate notifications for:

- new comments
 - replies to their comments
 - whenever someone else tags their username
 - when users apply peer feedback ratings to any of one's content
- It's also advisable that communities should be able to offer flexible *options* for users to *send or not send* notifications to those *users who wish to receive them*, for at least *some* types of signals.
 - * for example: don't notify someone of your emoticon "response" to their comment unless (1) they choose to receive that notification-type and (2) you want them to in this specific instance

Multimedia responses and live/ streamed channels

It should be fundamentally possible and easy for participants in media networks to:

- share rich recorded media messages (such as voice and video) whenever mutually desired
- initiate live/ streaming communications whenever mutually desired
- Attach rich media items and live channels directly to relevant media items

See [Meetings and Messages](#) regarding the relationship of meetings and messages in Structured Conversation.