

Trust-Based Media Flows for the Fediverse

Research & Development Focus

This project places research and development at its core, exploring how trust-based moderation and content flows can function as a viable alternative to traditional, centralised control systems.

Current models rely on top-down moderation, where authority is both concentrated and enforced. In contrast, this project investigates how trust networks - built through relationships, transparency, and shared practice - can enable moderation through bottom-up, community-driven flows.

This work leverages the established commons-based infrastructure of the Fediverse, built on ActivityPub, and extends it toward more coherent, flow-based news and media pathways.

Note: We frequently use the term “moderation.” While this relates to moderation tools available on Fediverse platforms (Mastodon, PeerTube, etc.), here it primarily refers to the flows of data moving between instances.

What is Being Built

This project will deliver a working [proof-of-concept codebase](#) that implements trust-based content and moderation flows across existing Fediverse platforms. Rather than building yet another new platform, we are developing a standards-based “flow layer” that interoperates with existing implementations.

Initially, we will focus on establishing compatibility with:

- Mastodon
- PeerTube
- Lemmy (or another federated forum implementation)

The key innovation is the development of **managed content flows between codebases**, implemented as working software:

- A middleware/service layer facilitating content movement, based on levels of trust between federated instances
 - A trust-based moderation model, where visibility and propagation emerge from these relationships
 - A reference implementation demonstrating these flows in practice
-

Why This Does Not Already Exist

Attempts to moderate content across the Fediverse are often hierarchical and per-instance, with each instance acting as its own “kingdom.” Many approaches mimic commercial social media, replicating familiar moderation patterns while overlooking systemic flaws.

Generations raised to think hierarchically - combined with programmers’ preference for control and predictable systems - have steered solutions away from more organic, communal approaches. As communities grow, top-down moderation becomes increasingly burdensome: moderators may experience trauma from exposure to harmful content, while also risking unintentional censorship. This highlights the need for trust-based, horizontal moderation flows that can scale without replicating the failures of commercial platforms.

While the Fediverse enables federation, it currently lacks:

- Coherent cross-platform content flows
- Trust-based propagation mechanisms beyond chronological or algorithmic feeds
- A shared commons model for native news distribution

This project addresses these gaps by connecting existing systems into a functional news/media ecology, rather than creating new silos.

Note: “Cross-platform” here refers specifically to different Fediverse platforms (e.g. Mastodon, PeerTube).

R&D Hypothesis

The central research question is:

Can trust-based moderation and distribution flows replace algorithmic amplification in a federated news ecosystem?

This will be explored through:

- Designing and implementing trust-driven flow logic
 - Testing propagation across trusted vs. untrusted nodes
 - Evaluating usability, resilience, and moderation outcomes
-

Mapping the Current Landscape

A core R&D task is to document and analyse existing Fediverse tools and flows, including:

- PeerTube (video distribution and moderation flows)
- Lemmy and similar forum systems (aggregation and ranking)
- Mastodon (real-time social flows)

This work will:

- Identify gaps between current implementations and the needs of trustworthy, native news production
- Compare flow-based models with legacy broadcast and platform approaches

The result will be a shared knowledge base of what works, what fails, and what can be recombined.

Trust-Based Flow Architecture

The technical foundation is a flow-based design:

- Content moves through networks based on trust relationships, not algorithmic amplification
- Moderation emerges from contextual trust, not universal rules
- Flows slow or stop where trust is weak (#KISS principle)

This produces systems that are:

- **Resilient** – no single point of failure
 - **Trustworthy** – grounded in human relationships
 - **Sustainable** – socially maintainable over time
-

Advancing Horizontal Development

The project develops and documents best practices in commons-based development, including:

- Open processes (#4opens)
- Transparent decision-making
- Lightweight, non-bureaucratic governance

Outputs include:

- Reusable design patterns for trust-based systems
 - Documentation of successes and failures
 - Feedback into the wider Fediverse ecosystem
-

Concrete Deliverables & Milestones

Month 3

- Technical specification of the flow architecture
- Prototype flow service (routing between two instances)
- Documentation of existing Fediverse flow patterns
- Early integration with one platform

Month 6

- Cross-platform prototype connecting at least two systems (e.g. Mastodon and PeerTube)
 - Demonstration of trust-based moderation flows
 - Public code repository and documentation
 - User-facing prototype: **#makinghistory** test environment
-

Infrastructure Strategy

A hybrid infrastructure model combines:

1. Horizontal Scaling (Grassroots)

- Many small, low-cost, community-run instances
- Built-in redundancy for text and media
- Distributed resilience

2. Institutional Anchoring

Larger instances hosted by:

- Universities
- Libraries
- Public service organisations
- European institutions

This ensures:

- Stability without centralisation
 - Public-interest support for core infrastructure
 - Independence at the network edges
-

Sustainability Model

Sustainability is approached as a social and infrastructural challenge:

- Core infrastructure supported by institutional hosting
- Network resilience through distribution and redundancy
- Development sustained through commons collaboration, not extractive funding

This aligns with public-interest values and avoids dependence on advertising or venture capital.

European Dimension & NGI Alignment

The project directly supports European priorities:

- Digital sovereignty through decentralised infrastructure
- Development of public-service media alternatives
- Alignment with openness, interoperability, and democratic participation

It aligns strongly with Next Generation Internet goals:

- Human-centric internet development
 - Open, interoperable technologies
 - Trust and resilience as core principles
-

Accessibility Commitment

All outputs will follow WCAG 2.1 accessibility standards, including:

- Semantic HTML and screen reader compatibility
- Full keyboard navigation
- Dual-layer UX (simple and advanced modes)

Accessibility is treated as a core design requirement.

Work Packages

WP1: Research & Specification

- Architecture and system mapping

- Output: technical design documentation

WP2: Core Development

- Flow service and ActivityPub integration
- Output: working prototype codebase

WP3: UX & Prototype

- #makinghistory interface
- Output: testable user prototype

WP4: Testing & Documentation

- Community testing and iteration
 - Output: public documentation and reports
-

Sustainability Beyond the Grant

The long-term model combines:

- Grassroots distributed hosting
- Institutional support for core nodes and codeing updates.

Discussions with institutional partners are underway, with the goal of securing formal support.

Summary

This project moves beyond #openweb theory to deliver:

- A working [prototype](#) of trust-based Fediverse flows
- A clear R&D investigation into alternatives to algorithmic media
- A practical, user-facing demonstration
- Reusable knowledge and tools for the wider ecosystem

The first-stage grant focuses on proof of concept and validation, followed by a bigger second-stage proposal to deliver a fully production-ready system.
