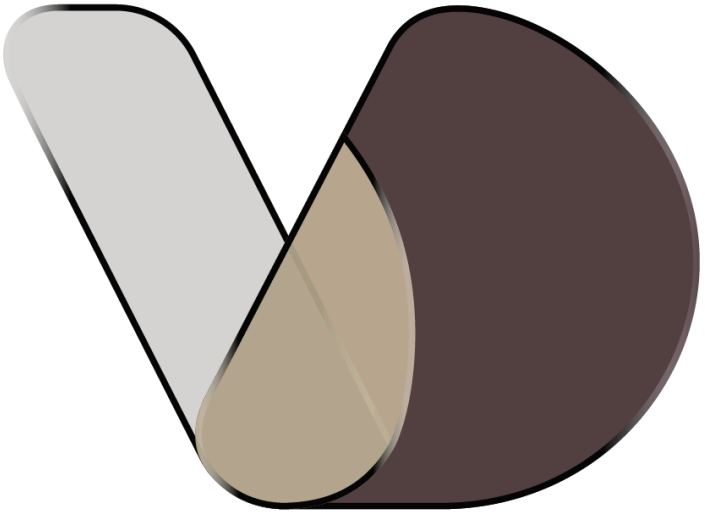

Vidi

Internet's living room

Vidi is the social media **people need**
This is how Europe can **make it a reality**

Pre-publication (17/04/2026)

Executive Summary

Vidi is a next-generation European video platform designed as a public–private digital infrastructure, **jointly operated by public broadcasters**. It aims to establish a trusted, democratic digital public space while enabling innovation and sustainable business models. Crucially, Vidi is not only conceived as a sovereign alternative to existing platforms, but also as a **compelling service** that users and creators will actively choose.

Part I - Why Vidi should exist

Vidi addresses the structural weaknesses of today's social media ecosystem: opaque governance, extractive business models, and weak accountability.

A sustainable and ethical business model: Unlike platforms reliant on surveillance-driven advertising, Vidi combines public infrastructure, public subsidies, and complementary private revenues (subscriptions and privacy-preserving ads). This hybrid approach ensures financial sustainability while aligning incentives with the public interest. Additional mechanisms such as peering fees ensure fair cost-sharing across the ecosystem.

A decentralized and sovereign infrastructure: Built on a decentralized protocol (based on the AT Protocol), Vidi enables multiple public broadcasters to operate interoperable infrastructure components. This architecture enhances resilience, scalability, and European digital sovereignty while remaining open to private innovation through APIs and third-party services.

Democratic governance and accountability: Governance is shared among public broadcasters, ensuring editorial independence and resistance to political capture. A Creators' Council complements this structure, giving content creators a formal role in decision-making. As a publicly governed platform, Vidi operates under democratic oversight and full compliance with European regulation.

Trust through transparency: Vidi is open-source and designed to be auditable by citizens, researchers, and institutions. It actively supports independent research and provides meaningful transparency into algorithms, moderation, and platform dynamics; contrasting with the opacity of dominant platforms.

A safer and more responsible ecosystem: Content moderation is decentralized and aligned with national legal frameworks, while strong safeguards protect minors and vulnerable users. Vidi integrates fact-checking partnerships, academic research, and community-driven tools to combat misinformation and harmful content at scale.

Part II - Why people will choose Vidi

Beyond its structural advantages, Vidi is designed to be an engaging, user-centric platform that redefines the social media experience.

Control over content and algorithms: Vidi introduces "algorithmic pluralism," allowing users to choose how content is recommended—whether through chronological feeds, curated selections, or alternative algorithms designed to reduce addictive patterns. This restores user agency over content discovery.

Higher-quality and more authentic content: The platform actively limits low-quality AI-generated content through identity verification for creators, rate limits, and incentive structures that favor human creativity. Partnerships with public broadcasters ensure a strong foundation of high-quality, trustworthy content from day one.

A better experience for creators: Vidi offers transparent monetization, fairer revenue-sharing models, and predictable income streams supported by public funding. Tools for cross-posting and gradual migration reduce the risk for creators transitioning from existing platforms.

Safer and more meaningful environments: Safe default settings, dedicated youth experiences, and community-specific moderation create spaces tailored to different audiences. Users can engage in environments aligned with their preferences without being exposed to unwanted or harmful content.

A social, community-first platform: Vidi reimagines online video as a shared experience rather than a solitary activity. Features such as watch parties, community spaces, and licensed reaction streams foster real-time interaction and collective engagement; positioning Vidi as the "Internet's living room."

Adoption and roadmap

Vidi's growth strategy combines institutional support and product attractiveness. Public broadcasters provide immediate access to high-quality content, while independent creators are onboarded progressively through strong incentives and monetization opportunities. Interoperability and cross-platform tools lower adoption barriers, while stricter enforcement of European digital regulations will further support migration over time.

In essence, Vidi is both **a systemic alternative and a desirable product**: a platform that not only fixes the structural flaws of social media but also offers a richer, safer, and more human-centered user experience.

Table of Contents

Context

With the mass adoption of the Internet, citizens now live a significant part of their social, cultural, and political lives online. Our public space no longer ends at the borders of our cities and institutions; it has expanded into a digital public space. Yet over the past two decades, this space has gradually been captured by a small number of Big Tech companies, which have transformed it into private property.

These platforms exercise unilateral control over what can be said, seen, and amplified online. Their economic incentives have led them to design addictive and polarizing features that maximize "engagement" and profit, too often at the expense of the public interest. The widespread use of algorithmic feeds optimized for attention has contributed to toxic online environments, undermined young people's mental health, and exacerbated social and political divisions.

Regulation alone (even in its most ambitious forms) has so far had only a limited impact. As long as citizens are given no credible alternative, societies remain effectively locked into privately owned digital public spaces governed by rules they did not choose. Today, the choice remains between accepting these toxic environments out of spite or withdrawing from public digital life altogether.

This report explains how Europe can reclaim its digital space with the creation of an attractive video-based social media built on top of public broadcasters.

Introduction

Vidi will be a European video platform jointly operated by public broadcasters. It is designed as a public-private digital infrastructure that ensures a safe and trusted digital public space, while enabling private actors to develop sustainable services and business models.

The platform builds on existing decentralized social media protocols developed through earlier experiments. It is fully interoperable, allowing private actors within Europe and public or private actors outside Europe to connect to Vidi. This interoperability enables users across different services to exchange content and interact with European citizens present on the platform.

Governance is shared among participating public broadcasters to ensure institutional stability, editorial independence, and the protection of freedom of expression. Content moderation is decentralized across European states, so that users access content in compliance with their local legislation.

Vidi also offers a plurality of content recommendation systems, giving users a choice between different experiences. Unlike existing video platforms designed primarily to maximize engagement, Vidi is built to promote authentic, safe, and meaningful interactions.

This document presents a detailed overview of the core components of Vidi and outlines a roadmap for its development and the progressive migration of users and creators.

Vidi will not be a "European Youtube/TikTok", it is meant to be a digital space with unique features. Vidi will transform video watching into a collective activity instead of an isolating exercise. Vidi will be the digital couch on which you sit and watch content with your friends and family.

Part I

Why you should

be on Vidi

Chapter I

Business model: sustainable and aligned with the public interest

The business model is the main blind spot of most existing alternative social media platforms. The absence of a clear and sustainable funding strategy fuels skepticism among potential users and institutional partners. More importantly, it creates a structural risk: when revenues are uncertain, platforms are eventually pushed toward intrusive advertising[1] and data exploitation to survive.

This concern already applies to several emerging platforms. For example, Bluesky currently relies on venture capital funding, which is being gradually consumed. While the project has proposed potential revenue models[2], none have yet been demonstrated at scale, raising concerns among prospective users[3]. Vidi is designed to avoid this trap. From the outset, it adopts a dedicated business model that ensures financial sustainability without relying on surveillance capitalism.

To achieve this, Vidi draws **inspiration from traditional European media** systems and combines three complementary funding pillars: a public infrastructure, public subsidies, and private funding sources through advertising and subscriptions.

A. Public infrastructure

The business model of Vidi rests on a **stable public infrastructure** provided by state actors. This infrastructure supports essential functions such as data storage and baseline content moderation. It guarantees that user data is hosted in compliance with European law and that illegal content is consistently removed.

This public infrastructure exposes open APIs that allow private actors to build and operate services on top of it. Infrastructure responsibilities are decentralized across Europe, with each participating country managing its share of storage and legal moderation.

This model can be compared to the European railway system. States maintain and interconnect the tracks, while private operators use this shared infrastructure to provide services and

[1] Rohrer, Christian, and John Boyd. "The rise of intrusive online advertising and the response of user experience research at Yahoo!." *CHI'04 extended Abstracts on human factors in Computing systems*. 2004.
[2] https://bsky.social/about/blog/7-05-2023-business-plan
[3] https://johnbattelle.medium.com/will-bluesky-embrace-advertising-yes-but-the-question-is-how-002b537a260a

compete. Similarly, Vidi provides a common digital backbone on which diverse actors can innovate.

Naturally, this public infrastructure does not prevent private actors from deploying their own technical stacks. However, recent history shows that privately owned infrastructures often adopt poor practices in moderation[4] and data protection[5,6]. When challenged, dominant platforms have repeatedly held users hostage[7] by offering a single choice: accept their problematic infrastructure or leave the digital public space altogether. By providing a public alternative, Vidi breaks this dependency and guarantees European citizens access to at least one trustworthy digital infrastructure.

B. Public subsidies

Operating a large-scale social media service, like running a traditional media outlet, is costly. At the same time, users have become accustomed to free access. Existing platforms have largely covered these costs through advertising based on personal data, a model that incentivizes toxic practices such as attention manipulation and addictive design[8].

To break this dynamic, Vidi introduces alternative funding sources for private actors. Drawing on established media funding mechanisms, **public subsidies** would be allocated based on measurable criteria such as audience size. These subsidies are not intended to cover all operating costs, but to provide a stable and predictable revenue base.

In return, subsidies would be conditional on compliance with European legislation, including rules on content moderation, competition, and data protection. This approach aligns economic incentives with public interest objectives.

C. Advertising and subscriptions

Public subsidies alone are not intended to fully fund social media services. Private actors remain free to **complement their revenues** through advertising and subscriptions.

While advertising has become problematic on existing platforms, the presence of public subsidies reduces the pressure to rely on highly intrusive and extremely profitable ad technologies. This creates incentives to favor privacy-preserving advertising models, even if they generate lower margins.

---

[4] https://www.npr.org/sections/alltechconsidered/2016/11/17/495827410/from-hate-speech-to-fake-news-the-content-crisis-facing-mark-zuckerberg
[5] https://www.reuters.com/sustainability/boards-policy-regulation/google-pay-135-million-settle-android-data-transfer-lawsuit-2026-01-28/
[6] https://fr.euronews.com/my-europe/2025/09/04/cookies-et-vie-privee-google-et-shein-sanctionnes-par-la-cnil-en-france
[7] https://www.dw.com/en/google-threatens-australia-pullout-in-row-over-news-content/audio-56313582
[8] Montag, Christian, and Jon D. Elhai. "On social media design,(online-) time well-spent and addictive behaviors in the age of surveillance capitalism." *Current Addiction Reports* 10.3 (2023): 610-616.

Combined with strict enforcement of GDPR and other European digital regulations, these three funding sources enable the emergence of a public–private video social media ecosystem built on a healthy, sustainable, and non-extractive business model.

D. Peering fees

As described in the protocol section, Vidi will rely on a decentralized architecture that allows private actors to build applications and services on top of the network. Private companies may therefore benefit from content that has already been moderated, indexed, and curated by public broadcasters. While this interoperability encourages innovation, it also raises the question of fair cost allocation. Content moderation, storage, and indexing are resource-intensive processes, and the actors financing these public-interest functions must not bear the burden alone if third parties derive commercial value from them.

To address this, Vidi can draw **inspiration from the telecommunications sector**, where operators occupy different positions within the network hierarchy. In the Internet ecosystem, Tier 1 operators manage backbone infrastructure and exchange large volumes of traffic, while smaller operators pay peering or transit fees[9] when relying on this infrastructure. A comparable model could apply within Vidi: public broadcasters operating core infrastructure could charge usage-based fees to private actors who depend on their infrastructure for data access and content discovery.

These peering fees would be proportional to traffic volume or resource consumption, ensuring that costs are distributed fairly while preserving openness. Such a mechanism would sustain the public infrastructure, prevent free-riding, and create a balanced ecosystem where private innovation coexists with publicly financed digital foundations.

[9] https://en.wikipedia.org/wiki/Peering

Chapter II

Protocol: building on existing decentralized ecosystems

Operating a jointly governed social media platform is structurally complex. Each stakeholder has distinct resources, constraints, and priorities. Because Vidi relies on public broadcasters across multiple countries, each participant must retain a degree of operational autonomy. At the same time, the platform should remain open enough for private actors to build services and businesses on top of it.

Rather than relying solely on complex legal arrangements, decentralization can be embedded directly into the technical architecture. A decentralized protocol makes cooperation more scalable, distributes responsibilities, and reduces governance friction. This section outlines the protocol choices that enable Vidi to be jointly operated by public broadcasters while remaining interoperable with external actors.

Decentralized social media is not a new field. Over the past decade, multiple projects[10] such as Mastodon[11], Bluesky[12], and PeerTube[13] have explored different architectural models. Vidi can build on these experiments, reuse proven components, and avoid known scalability and governance limitations[14,15].

A. The need for decentralization

A straightforward approach would be to create a single public company jointly owned by participating broadcasters and operate a centralized platform. While simpler to launch, this model creates long-term structural problems.

First, participation will likely be progressive. Some countries may join later after observing initial results. Integrating new public stakeholders into a pre-existing multinational company is politically and administratively heavy. Second, unequal financial contributions would create recurring tensions over control, priorities, and resource allocation.

---

[10] Damie, Marc, and Edwige Cyffers. "Fedivertex: a Graph Dataset based on Decentralized Social Networks for Trustworthy Machine Learning." *Proceedings of the ACM Web Conference*. 2026.
[11] https://joinmastodon.org/
[12] Kleppmann, Martin, et al. "Bluesky and the at protocol: Usable decentralized social media." *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*. 2024.
[13] https://joinpeertube.org/
[14] Raman, Aravindh, et al. "Challenges in the decentralised web: The mastodon case." *Proceedings of the internet measurement conference*. 2019.
[15] Anaobi, Ishaku Hassan, et al. "Will admins cope? Decentralized moderation in the fediverse." *Proceedings of the ACM Web Conference 2023*. 2023.

A **decentralized infrastructure** addresses these issues more effectively. Each public broadcaster can deploy and operate its own part of the network while remaining connected through a shared protocol. New participants (possibly outside Europe!) can join at any time by deploying compatible infrastructure, without requiring institutional restructuring. Funding and operational responsibilities remain clearly partitioned.

This architecture also improves **accountability**. If moderation or service quality is insufficient for content originating from a given region, responsibility can be traced to the corresponding operator. The protocol does not eliminate governance challenges, but it makes them more tractable by aligning technical boundaries with institutional responsibilities.

The logic is comparable to the European Union itself: cooperation through shared rules and interoperable systems rather than absorption into a single centralized state. Decentralization creates a balance between sovereignty and coordination, enabling a **shared public space with distributed control**.

Finally, a decentralized infrastructure increases **resilience to political pressure**. There is no single technical or legal control point that can be captured to dominate the platform. As with other decentralized networks, this structure strengthens institutional robustness and protects operational independence.

B. Existing decentralized social media

Vidi does not need to reinvent decentralization from scratch. Several protocols for decentralized social media already exist, and the objective is to learn from these experiments and reuse their strongest design patterns. Two protocols currently structure this space: ActivityPub and the AT Protocol.

ActivityPub[16] is a standardized protocol used by platforms such as Mastodon and PeerTube. It creates a federation of servers where each server manages its own users, data, and moderation rules. Interoperability is achieved because users on one server can follow accounts and access content hosted on another. ActivityPub has gained strong adoption within the open-source ecosystem, and the broader ActivityPub network, known as the Fediverse, counts millions of users.

However, ActivityPub presents two structural limitations at a very large scale. First, its server-centric model creates user friction during onboarding[17]. New users are often confused by the need to choose a server and understand federation rules, which complicates adoption beyond technically literate audiences. Second, the model duplicates operational responsibilities across servers. Each server must implement its own discovery, storage, and moderation processes. The same content may be reviewed and processed many times across the network, with limited native support for shared moderation layers[18,19]. This redundancy increases operational costs and complicates scaling to hundreds of millions of users.

---

[16] Prodromou, Evan. *ActivityPub*. " O'Reilly Media, Inc.", 2024.
[17] https://www.theverge.com/tech/881352/mastodon-default-server-recommendations-experiment
[18] Agarwal, Vibhor, et al. "Decentralised moderation for interoperable social networks: A conversation-based approach for pleroma and the fediverse." *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 18. 2024.
[19] https://writings.thisismissem.social/open-source-tools-for-the-future-of-decentralized-moderation/

The AT Protocol[20] is more recent. It was designed by former Twitter engineers and serves as the foundation for Bluesky. It was designed with two primary goals: large-scale performance and resilience to political or institutional pressure.

Unlike ActivityPub, the AT Protocol separates infrastructure roles into specialized components. It defines three main server types: Personal Data Servers (PDS), Relays, and Feed Generators:.

- *Personal Data Servers* store user data and handle account-level interactions.
- *Relays* aggregate and index content across the network and perform resource-intensive discovery tasks.
- *Feed Generators* consume relay data and produce customized feeds.

This pipeline architecture enables resource pooling: many feed generators can depend on the same relay infrastructure, which improves efficiency and scalability. Decentralization can then occur at different layers, with varying cost and responsibility.

This modular structure is more technically complex than ActivityPub, but it is better suited for very large user bases. It also enables a smoother user experience because account identity and server location are less visible in everyday usage. Bluesky has demonstrated that users can migrate from centralized platforms with relatively low onboarding friction.

The AT Protocol is not without limitations. In practice, its main deployment today is Bluesky, where most infrastructure remains concentrated under a single company. Decentralization is supported by the protocol design but not yet widely realized in operations. By contrast, ActivityPub already runs across tens of thousands of independently operated servers. The AT Protocol therefore appears to be a strong technical foundation, but it must be adapted and deployed in a genuinely distributed way to meet Vidi's governance and sovereignty objectives.

C. Adapting the AT Protocol to Vidi

Vidi will use the AT Protocol as its primary technical foundation. The **core infrastructure** will be **deployed by participating public broadcasters**. Each broadcaster joining the network will operate its own Personal Data Servers, Relay servers, and Feed Generators. From the user perspective, however, this decentralized infrastructure will be unified behind a single, simple Vidi application interface. The complexity remains at the protocol layer, not at the user interface layer.

It is important to stress that Vidi is built on an **open protocol**. Any private actor can deploy their own Feed Generators, Relays, Personal Data Servers, or client applications to build services and businesses on top of the network. Public broadcasters provide a trusted core infrastructure, but they do not hold a technical monopoly. This openness enables innovation, competition, and a broader interoperable ecosystem around Vidi.

To **optimize operations** and moderation quality, responsibilities will initially be spread. For example, German content would be handled by the German broadcaster's infrastructure, while

[20] Kleppmann, Martin, et al. "Bluesky and the at protocol: Usable decentralized social media." *Proceedings of the ACM Conext-2024 Workshop on the Decentralization of the Internet*. 2024.

other relays would forward such content to the appropriate language relay. This model provides two operational benefits: content is reviewed by actors with the relevant linguistic and cultural expertise, and moderation workload is distributed proportionally rather than concentrated in larger countries. Assuming relatively uniform content production per user across Europe, operational costs would scale roughly with population size.

This country-based allocation is a practical starting point but not the long-term optimal model. Empirical research on online behavior shows that social interactions cluster strongly by language across national borders. A language-based relay structure, such as a shared French-speaking or Spanish-speaking relay, would better reflect real usage patterns. Such cross-border operational units would require additional institutional cooperation and are beyond the scope of this concept note, but the protocol design should make them technically straightforward to implement.

The AT Protocol should be considered a base layer rather than a fixed framework. Targeted extensions will be required to support Vidi's operational model. In its current form, the protocol does not emphasize relay-to-relay coordination. For Vidi, horizontal communication between relays would improve routing efficiency and moderation workflows, for example when content is first received by a relay outside its primary language domain. Adding structured inter-relay signaling and workload routing would strengthen distributed operations.

In addition, the AT Protocol was initially designed around microblogging use cases and will require adaptation for video-heavy workloads. Video platforms introduce specific constraints in storage, streaming, transcoding, and delivery. Vidi will therefore also draw from design patterns and components developed in ActivityPub-based video systems, particularly PeerTube. Peer-to-peer assisted video delivery, for example, can significantly reduce server load and bandwidth costs at scale. Reusing and integrating such mechanisms will be essential to keep infrastructure costs under control while maintaining performance.

Vidi's protocol strategy is therefore pragmatic: adopt the AT Protocol as the structural backbone, extend it where necessary for governance and moderation coordination, and incorporate proven video-scale components from other decentralized platforms.

Chapter III

Governance: decentralized and democratic

As a public good, Vidi is designed with governance fundamentally different from existing social media platforms. For the first time, a social media platform can become a truly democratic digital space.

A. Federated public broadcasters

The infrastructure underlying Vidi will be jointly operated by European public broadcasters, placing them at the core of the platform's governance. These broadcasters will collectively define key policies, including content moderation rules. Anchoring the platform in such a federation provides a stable foundation for private actors building services on top of Vidi.

Policies will not depend on a single country. If undemocratic political powers emerge in one European state, they cannot unilaterally change platform rules. Policy changes require consensus across participating broadcasters and their respective governments, ensuring stability and preventing politicization of the platform.

As established public actors, broadcasters are **accountable to their citizens** and to European values. Governance through broadcasters should therefore be understood as a form of **indirect democracy**, reflecting the interests and culture of local users while maintaining a European-wide framework.

To formalize this collaboration, a limited EU-wide legal structure may be required to centralize specific functions, such as shared tool development, coordination, or intellectual property management. This entity should exist solely to support the platform's continuity and coherence, not to operate or govern it. Operational control and policy decisions must remain with the public broadcasters within the federation. The central structure should therefore be intentionally lightweight, preserving decentralization as a core principle. This design enhances resilience to political pressure and facilitates the gradual integration of new members, allowing Vidi to begin with a small group of founding countries while remaining open to future expansion.

B. Creators' council

Content creators are the backbone of any social media platform, yet private platforms have historically ignored their perspectives, making unilateral decisions that directly impact creators' livelihoods. The global, dispersed nature of creators makes organizing them collectively challenging, leaving them with limited influence and high revenue instability.

Vidi addresses this gap by establishing a Creators' Council to represent the **interests of content creators**. This council will participate in decision-making processes, following a model similar to Wikipedia's collaborative governance[21,22], which successfully coordinates contributors across cultures and languages.

Together, the federation of public broadcasters and the Creators' Council form the two main governance bodies of Vidi. This structure balances the interests of infrastructure providers and content creators, while creating a resilient shield against political interference within individual broadcasters.

C. Interconnected social media

Vidi is an open, interoperable platform. **Private companies can build services** and products on top of or interconnected with Vidi. While no separate governance body for these actors is proposed at launch, their perspectives may be incorporated over time, particularly if European businesses develop sustainable services reliant on the platform. This ensures that future decisions can account for the operational needs of third-party services without compromising the platform's public-interest mission.

[21] Forte, Andrea, Vanesa Larco, and Amy Bruckman. "Decentralization in Wikipedia governance." *Journal of Management Information Systems* 26.1 (2009): 49-72.
[22] Konieczny, Piotr. "Adhocratic governance in the Internet age: A case of Wikipedia." *Journal of information technology & politics* 7.4 (2010): 263-283.

Chapter IV

Content curation: safe and personalized content

As social media platforms scale, the volume of unsafe and harmful content inevitably increases. While European legislation strongly incentivizes effective content moderation[23], profit-driven platforms have consistently under-invested in moderation capabilities[24], prioritizing growth and revenue optimization over user safety.

Vidi follows a different logic. As a publicly owned digital infrastructure, its governance is not guided by profit maximization. Ensuring user safety is a core public service mission of the institutions operating the platform.

Content curation within Vidi is therefore structured across multiple layers and distributed between Relays and Feed Generators. This architecture enables robust enforcement of legal requirements while allowing flexibility and personalization at the user level.

A. Illegal content moderation

The most critical layer of content curation concerns the detection and removal of illegal content, such as child sexual abuse material. Within Vidi, this responsibility primarily lies with relay servers, which act as gatekeepers before content is distributed through feeds.

Data storage servers may also contribute to this effort by detecting illegal content at upload time and preventing its propagation across the network. As storage providers may have limited computational resources, such preventive detection is not mandatory. However, storage servers remain accountable for the content they host. Relays may restrict or block servers that repeatedly fail to meet legal and operational standards, based on an accumulated reputation. In this effort, GIFCT[25] and Stop NCII[26] databases will be integrated to each of the relevant layers to prevent the diffusion of flagged materials.

In the Vidi architecture, each participating public broadcaster operates a **national relay server**. Moderation responsibilities are therefore distributed along linguistic and cultural lines. Each relay focuses on content produced in its national language or languages, ensuring that moderation decisions are made by institutions with the necessary linguistic competence and cultural context. This approach addresses a well-documented limitation of major platforms, which often lack sufficient moderation capacity for less widely spoken European languages.

---

[23] Nunziato, Dawn Carla. "The Digital Services Act and the Brussels effect on platform content moderation." *Chi. J. Int'l L.* 24 (2023): 115.

[24] https://www.cnbc.com/2026/03/19/meta-cut-back-third-party-vendors-favor-of-ai-for-content-enforcement.html

[25] https://gifct.org/hsdb/

[26] https://stopncii.org/

Decentralized moderation across European countries also creates a structured framework for cooperation and alignment between national authorities. While moderation remains compliant with local legislation, this shared infrastructure encourages legal convergence and mutual oversight, reducing the risk of arbitrary or politically motivated content restrictions that could undermine fundamental rights.

B. Preferential content moderation

A common misconception about content moderation is that it ends with the removal of illegal material. In practice, platforms also moderate content based on editorial policies or community standards. Content such as nudity, depictions of violence, or firearms may be lawful, yet still inappropriate for certain audiences or contexts. This form of "preferential" moderation is essential to maintaining safe and inclusive digital spaces, particularly for minors and vulnerable groups.

Within Vidi, preferential moderation is primarily handled by feed generators. Each feed is **designed for a specific audience** or community and defines its own moderation policies accordingly. Feed generators are therefore responsible for ensuring that the content they distribute complies with their declared rules. This enables the creation of feeds tailored to specific needs, such as child-friendly environments, safe spaces for discriminated populations, or thematic feeds focused on particular interests, for example video games, culture, or education.

Rather than imposing a single, universal set of norms on all users, Vidi supports a pluralistic model of content curation. Where dominant platforms concentrate all users into a single, highly constrained public space, Vidi enables a constellation of smaller, purpose-driven digital spaces. This approach allows users to engage in environments aligned with their expectations and values, without being exposed to unwanted or harmful content.

Preferential moderation can be resource-intensive, particularly for video content. To reduce duplication of effort across the ecosystem, relay servers support feed generators by providing standardized content labels, such as indicators for nudity, violence, or sensitive themes. As relays already perform content analysis for illegal content detection, these additional classifications can be derived from the same processes.

This division of responsibilities is designed to minimize overall operational costs. By centralizing baseline content labeling at the relay level, Vidi avoids requiring every feed generator to independently analyze all incoming video content, resulting in a more efficient and scalable moderation system.

C. Pluralistic content recommendation

Once content has been moderated, a final and decisive step remains in the curation pipeline: content recommendation. Moderation defines what can be shown, but recommendation determines **what is shown**. Given the scale of video catalogs, effective recommendation mechanisms are essential to help users navigate content meaningfully.

Dominant platforms have centralized this function within opaque, proprietary algorithms. These systems are typically optimized for engagement metrics, such as watch time or interaction frequency, rather than for user well-being or informational quality. As a result, users are subjected to recommendation logic they neither understand nor control.

Vidi adopts a fundamentally different approach by promoting **algorithmic pluralism**[27,28]. Instead of enforcing a single, engagement-driven recommender system, the platform allows users to choose among multiple recommendation models, each based on distinct principles and objectives. These may include chronological feeds, interest-based recommendations, editorially curated selections, or algorithms explicitly designed to minimize addictive patterns and polarization.

By making recommendation systems **transparent and customizable**, Vidi restores user agency over how content is prioritized and consumed. Recommendation becomes a configurable public service rather than an invisible mechanism optimized for profit.

In addition to algorithmic approaches, Vidi actively reinstates human-driven recommendations. Trusted actors such as journalists, educators, cultural institutions, or content creators may curate "watch lists" highlighting content they consider relevant, informative, or appropriate for specific audiences. These curated lists can coexist with algorithmic feeds and offer an alternative path through the platform's content ecosystem.

Recommendation algorithms may be developed either by Vidi's maintainers or by independent third parties. These algorithms would be accessible through a public catalog from which users can select their preferred recommender system. Some providers may choose to offer premium algorithms, particularly where maintaining high-quality manual or computationally intensive curation induces significant costs. However, even in the case of paid algorithms, Vidi should impose baseline transparency requirements, enabling researchers and oversight bodies to detect systemic biases, manipulation risks, or harmful dynamics.

This hybrid model is particularly valuable for younger audiences. Children are especially vulnerable to the risks associated with engagement-optimized algorithms, including compulsive consumption and exposure to inappropriate material. For child-oriented experiences, Vidi can therefore prioritize manual or editorial recommendation models, ensuring a safer and more intentional content discovery process.

Through pluralistic and transparent recommendation mechanisms, Vidi aligns content discovery with European values of autonomy, diversity, and public accountability, while allowing users to shape their own digital experience.

[27] https://www.lemonde.fr/idees/article/2024/09/25/pour-le-pluralisme-algorithmique_6332830_3232.html

[28] Jain, S., et al. "Algorithmic pluralism: A structural approach to equal opportunity." *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 2024.

Chapter V

Youth protection: safe defaults and dedicated experiences

Multiple scandals have revealed the harmful effects that major social media platforms can have on young users[29,30,31,32,33,34,35,36]. As a result, the safety of minors online has become a major topic in policy debates[37]. While there is broad agreement that action is necessary, many currently proposed solutions remain unsatisfactory. On the one hand, large platforms often attempt to evade responsibility. On the other hand, some policymakers advocate for outright social media bans for minors; an approach that many experts consider ineffective and potentially harmful. In this context, Vidi offers an alternative perspective and proposes a third path.

A. Safe default settings

The first essential building block of digital safety is the establishment of safe default settings. The **default experience** should be **appropriate for all** users, including minors. In practice, this means that certain sensitive topics (for example, graphic war-related content) would not be visible by default. Access to such material would require users to verify that they are above a specified age threshold.

More broadly, any feature restricted for minors under European or national legislation would be disabled by default. These safe defaults may also benefit adult users who prefer a more filtered social media experience.

Adopting this approach also acknowledges the legitimate role of sensitive content in the public sphere. For example, content related to ongoing conflicts should not be accessible to children, but it remains important that adults can access such information to stay informed.

Current major platforms often apply inconsistent moderation policies: certain topics such as abortion or war may be censored or suppressed, while other potentially harmful content (such as material promoting gun ownership) remains widely promoted. A safe-default approach would

[29] https://www.wsj.com/articles/the-facebook-files-11631713039
[30] https://www.amnesty.org/en/documents/POL40/7350/2023/en/
[31] https://www.amnesty.org/en/documents/POL40/7349/2023/en/
[32] https://nypost.com/2022/10/08/molly-russells-death-may-spark-social-media-reforms-for-us-kids/
[33] Al-Samarraie, Hosam, et al. "Young users' social media addiction: causes, consequences and preventions." *Information Technology & People* 35.7 (2022): 2314-2343.
[34] https://www.bbc.com/news/articles/cze3p1j710ko
[35] https://www.nytimes.com/2017/11/04/business/media/youtube-kids-paw-patrol.html
[36] https://www.wired.com/story/youtube-minecraft-among-us-disturbing-videos/
[37] Ingber, Alexis Shore, and Chris Chao Su. "Protecting children of the TikTok era: A discourse analysis in the absence of law." *Policy & Internet* 17.2 (2025): e431.

instead allow the emergence of a digital space that protects vulnerable audiences without distorting public debate.

B. Dedicated feeds

Beyond safe defaults, it is also important to design **age-appropriate experiences** for children[38][39]. Platforms such as YouTube Kids demonstrate that many parents value environments specifically tailored to younger users.

However, safe defaults alone only ensure a child-safe experience, not necessarily a child-friendly one. Vidi will therefore develop dedicated feeds designed to be interesting and beneficial for younger audiences. Because Vidi will be operated by public broadcasters, the educational missions traditionally associated with these institutions can naturally extend to the platform.

Unlike platforms such as YouTube Kids, Vidi will not rely on engagement-maximizing algorithms designed to capture children's attention. Instead, it will provide curated feeds that contribute positively to the education and development of younger users.

To achieve this, collaboration with educators and child-development experts will be essential. The objective is to create experiences that are carefully designed for children, while avoiding the pervasive influence of algorithmic recommendation in their digital lives. Limiting algorithmic exposure may also reduce risks associated with recommendation systems, whose long-term effects are still not fully understood.

C. Supporting educators and parents

Creating a safe digital space for children is not purely a technical challenge; it also has an important educational dimension[40]. **Teaching young users** how to navigate the internet safely is widely recommended by experts in digital safety. Vidi will therefore actively support this educational mission.

One key component is helping parents guide their children online. This guidance typically involves parental control tools. Like most platforms, Vidi will provide such features. However, many parents are unaware of these tools or unsure how to use them effectively. Vidi will therefore also provide clear documentation and accessible training resources to ensure that as many families as possible can benefit from these protections.

In parallel, media literacy and digital education are increasingly integrated into school curricula across Europe. As a publicly operated platform, Vidi should collaborate with teachers and educational institutions to develop learning materials that explain topics such as algorithms, digital platforms, and online safety.

---

[38] Livingstone, Sonia, and Kim R. Sylwander. "Conceptualizing age-appropriate social media to support children's digital futures." *British Journal of Developmental Psychology* (2025).
[39] Livingstone, Sonia, and Kim R. Sylwander. "There is no right age! The search for age-appropriate ways to support children's digital lives and rights." *Journal of Children and Media* 19.1 (2025): 6-12.
[40] Goodyear, Victoria A., et al. "Approaches to children's smartphone and social media use must go beyond bans." *bmj* 388 (2025).

The objective of these collaborations would not be to promote Vidi itself, but rather to share the expertise of those operating the platform. Such cooperation could be particularly valuable for educators, as Vidi will offer a level of transparency rarely available on existing platforms. This transparency can help produce concrete and practical educational resources.

D. Age verification

Finally, age verification has recently become a prominent issue in European policy debates[41]. Many policymakers advocate for strict age verification mechanisms, and Vidi will naturally comply with any regulatory requirements adopted in this area.

However, it is important to acknowledge the significant **concerns raised by experts** regarding current technological approaches. For instance, hundreds of European security and privacy researchers have signed an open letter[42] calling for a moratorium on existing age-assessment technologies due to risks related to privacy, security, and reliability.

These concerns highlight the importance of carefully evaluating age verification mechanisms and prioritizing solutions that protect both minors and fundamental rights.

[41] https://www.dw.com/en/european-nations-mull-social-media-ban-for-children-instagram-tiktok-anxiety-depression/a-75841184
[42] https://csa-scientist-open-letter.org/ageverif-Feb2026

Chapter VI

Building trust: from transparency to accountability

Rebuilding a shared digital public space is a major challenge. Over the past decade, repeated scandals and the lack of transparency from major technology companies have significantly weakened public trust in social media platforms[43,44]. Establishing the trustworthiness of Vidi from the outset is therefore essential. This requirement is even more important given that the platform will rely partly on public funding for its operation. Vidi must therefore be exemplary in its governance and provide the transparency and accountability necessary to earn the trust of both policymakers and citizens.

A. Transparent architecture

Transparency will be the primary mechanism for building this trust. Vidi will be a decentralized and open-source platform, allowing anyone to inspect the underlying code and algorithms. In principle, **any citizen will be able to audit** how the platform operates.

Because such analysis requires advanced technical skills, Vidi's stakeholders should also invest in science communication and educational materials explaining the platform's architecture, algorithms, and governance in accessible terms.

This level of transparency contrasts sharply with current major platforms. Their executives often make strong claims[45,46,47,48,49] regarding the fairness, safety, and privacy of their systems, yet the public has no meaningful way to verify them. Repeated controversies have gradually eroded trust in these claims[50,51,52].

Vidi will take a different approach. Rather than asking the public to trust the platform, Vidi will prove that its systems are fair, safe, and privacy-preserving.

---

[43] https://www.gartner.com/en/newsroom/press-releases/2025-10-13-gartner-consumer-survey-identifies-most-trusted-social-media-platforms-for-information-accuracy

[44] https://news.gallup.com/opinion/gallup/357446/young-people-rely-social-media-don-trust.aspx

[45] https://www.wired.com/story/youtube-ceo-defends-efforts-reduce-violent-content/

[46] https://blog.google/innovation-and-ai/technology/safety-security/keeping-private-information-private/

[47] https://www.nytimes.com/2019/03/06/technology/mark-zuckerberg-facebook-privacy.html

[48] https://blog.x.com/en_us/topics/company/2023/an-update-on-our-work-to-tackle-child-sexual-exploitation-on-x

[49] https://blog.x.com/en_us/topics/company/2023/maintaining-the-safety-of-x-in-times-of-conflict

[50] https://www.nytimes.com/2024/06/22/technology/zuckerberg-instagram-child-safety-lawsuits.html

[51] https://www.theguardian.com/technology/2026/mar/05/mark-zuckerberg-meta-trial

[52] https://www.theverge.com/2024/10/1/24259653/snap-new-mexico-ag-lawsuit-csam-kids-safety

B. Open to researchers

The European Digital Services Act (DSA) requires large platforms to provide researchers with tools[53] enabling them to study social media activity and its societal impacts. In practice, however, some platforms offer only limited access to such data; or none at all[54].

Vidi will fully embrace this principle and actively develop tools that **facilitate independent research**. Researchers will be able to study the dynamics of the platform, its algorithms, and its societal effects under appropriate privacy safeguards.

As a **digital public space designed to connect citizens** rather than maximize profits, Vidi will welcome such investigations. Independent research can contribute to improving safety, understanding social dynamics online, and strengthening the overall resilience of the platform.

C. Accountable organization

Finally, Vidi will be built around strong institutional accountability. Large technology companies often avoid responsibility through lobbying or diplomatic pressure from the governments of the countries in which they are headquartered. Such dynamics can make democratic oversight difficult.

By contrast, Vidi will operate as a European public service. It will fully **comply with European legislation** and operate under democratic oversight. As a publicly funded initiative, it may also be subject to parliamentary scrutiny to ensure that public resources are used responsibly and that the platform operates in accordance with its mission.

In this sense, Vidi has the potential to become the **digital public service for the 21st century**. Achieving such an ambition requires more than technological innovation; it requires building an organization that consistently earns the trust of the citizens it connects.

[53] https://digital-strategy.ec.europa.eu/en/news/new-measures-unlock-access-data-largest-online-platforms-support-research
[54] https://www.techpolicy.press/what-the-x-fine-reveals-about-data-access-under-article-40-of-the-digital-services-act/

Part II

Why you want to be on Vidi

Chapter VII

Building a social media for and by humans

Major video platforms are increasingly **saturated with AI-generated content**[55,56,57] optimized to exploit recommendation algorithms and maximize revenue, often at the expense of authenticity and viewer satisfaction. This growing wave of low-quality automated content (sometimes referred to as "AI slop") is gradually degrading the quality of online public spaces. Many users are becoming frustrated with manipulative content, automated accounts, and algorithmic bait designed to capture attention rather than provide meaningful information or entertainment.

Vidi aims to address this problem directly. By design, the platform will implement mechanisms that limit the visibility and profitability of low-quality automated content. The objective is not to reject technological innovation, but to ensure that the digital public space remains **centered on human creativity**, cultural expression, and meaningful interactions.

A. New algorithms, new rules

The first mechanism for promoting human-created content will lie in the design of recommendation algorithms. A key advantage of algorithmic plurality is that it weakens the strategies commonly used by automated accounts and mass-produced AI content. Much of today's AI-generated media is optimized to exploit engagement-driven algorithms[58], which reward sensationalism, repetition, and volume.

Moving away from purely engagement-based systems significantly reduces the incentives that fuel the production of AI slop. In addition, the presence of manually curated feeds will shift attention toward high-quality content curated by trusted actors (such as journalists, educators, or creators) rather than content generated at scale by automated systems. This hybrid model helps restore visibility for authentic human creativity while **limiting the influence of algorithmic manipulation.**

---

[55] https://www.theguardian.com/global/commentisfree/2025/jan/08/ai-generated-slop-slowly-killing-internet-nobody-trying-to-stop-it

[56] https://theconversation.com/what-is-ai-slop-a-technologist-explains-this-new-and-largely-unwelcome-form-of-online-content-256554

[57] https://hls.harvard.edu/today/ai-content-and-algorithms-is-coming-for-your-kids/

[58] Hussain, Khalid, M. Laeeq Khan, and Aqdas Malik. "Exploring audience engagement with ChatGPT-related content on YouTube: Implications for content creators and AI tool developers." *Digital Business* 4.1 (2024): 100071.

B. Identity verification using EU Digital Wallets

Bot detection remains a major challenge on social media platforms[59]. Existing solutions typically face a difficult trade-off: they are either insufficiently effective or excessively intrusive. For example, CAPTCHA-based systems[60] are inexpensive and relatively privacy-preserving, but they are increasingly ineffective against sophisticated bots. Conversely, full identity verification can be highly effective but is costly to deploy at scale and raises significant privacy concerns.

Some recent social media initiatives aiming to combat automated accounts have proposed universal identity verification. However, applying such verification to the entire user base risks discouraging adoption and creating new security vulnerabilities. Storing sensitive identity data for millions of users would increase both operational costs and the potential impact of security breaches.

Vidi proposes a more **targeted approach focused on content creators**. On video platforms, creators represent a small fraction of users but play a disproportionately influential role in shaping the platform's content ecosystem. Concentrating identity verification on this group significantly reduces costs and avoids introducing unnecessary friction for users who simply want to watch or interact with content.

To **minimize security risks**, identity verification should rely on EU Digital Wallets[61,62,63], which are progressively being deployed across Europe. When implementing using proper cryptographic primitives[64,65,66], these systems provide secure and standardized identity verification while limiting the amount of personal data shared with platforms[67]. Importantly, verification would remain pseudonymous: Vidi would confirm that a creator is a verified individual without publicly revealing their identity.

This mechanism also supports legal **accountability.** If a content creator publishes illegal material, Vidi will be able to cooperate with authorities when required by law. As a publicly governed platform, Vidi must operate within clear legal responsibilities rather than attempting to avoid them.

By making identity verification mandatory for creators, Vidi significantly raises the cost of operating automated "AI slop" accounts. Creating large networks of bot creators would become both expensive and traceable, thereby disrupting one of the main drivers of low-quality automated content.

---

[59] Aljabri, Malak, et al. "Machine learning-based social media bot detection: a comprehensive literature review." *Social Network Analysis and Mining* 13.1 (2023): 20.
[60] Gossweiler, Rich, Maryam Kamvar, and Shumeet Baluja. "What's up CAPTCHA? A CAPTCHA based on image orientation." *Proceedings of the 18th international conference on World wide web*. 2009.
[61] https://ec.europa.eu/digital-building-blocks/sites/spaces/EUDIGITALIDENTITYWALLET/pages/694487738/EU+Digital+Identity+Wallet+Home
[62] Balani, Henry. "Securing the future of the EU Digital Identity Wallet: Why we need corporate digital identity standards now." *Journal of Financial Compliance* 9.1 (2025): 32-39.
[63] Vanella, Alessandro. *Evolution of Digital Identity in Europe: Experimenting with the eIDAS 2.0 Framework and the EU Digital Identity Wallet*. Diss. Politecnico di Torino, 2025.
[64] Podda, Emanuela, et al. "The impact of zero-knowledge proofs on data minimisation compliance of digital identity wallets." *Internet Policy Review* 14.3 (2025): 1-29.
[65] Ramos Fernández, Raül. "Regulatory options for integrating zero-knowledge proofs into the European Digital Identity Wallet." *International Review of Law, Computers & Technology* 39.3 (2025): 303-326.
[66] Saarela, Uuna. "Practical Zero-Knowledge within the European Digital Identity Framework: Implementing Privacy-Preserving Identity Checks." (2024).
[67] https://yivi.app/en/privacy_and_security/

C. Rate limitations

In addition to identity verification, Vidi will implement rate limits to further **disrupt the economic incentives** behind AI-generated content farms.

First, verified identities make it possible to limit the number of accounts that a single individual can control. These limits can be tightened in cases of repeated policy violations, such as the publication of illegal or abusive content.

Second, the platform will impose **default limits** on the rate at which creators can upload videos. Automated content farms typically publish very large volumes of content in short periods of time; behavior that differs significantly from typical human creators. By restricting upload frequency, Vidi can significantly reduce the profitability of such operations.

Exceptions to these limits could be granted when justified. For instance, media organizations or television channels that legitimately produce large quantities of content could request higher upload thresholds, subject to manual review.

Applying conservative limits by default (especially for newly created accounts) provides a powerful and practical tool to slow the spread of low-quality automated content while preserving flexibility for legitimate creators.

D. Collaborative labeling

As Vidi is designed as a platform serving internet citizens, moderation should not rely solely on a top-down approach. Users themselves can play an important role in maintaining the quality and integrity of the platform. Vidi will therefore implement mechanisms allowing the **community** to **flag problematic or low-quality content**.

When a clear consensus emerges among contributors, a public label can be attached to the content; similar to the community notes system used on X. These labels (e.g. AI generation) would inform viewers about potential issues and could also be taken into account by feed generators when deciding how prominently to recommend the content.

Importantly, this labeling process would be fully transparent. Labels would be visible to all users, and creators would have the opportunity to respond or contest them if necessary. This transparency helps avoid the unintended consequences associated with opaque moderation practices, such as shadow banning, which are common on many existing platforms.

Chapter VIII

Fake news: winning the information battle

Beyond the proliferation of AI-generated content, social media platforms are also deeply affected by another persistent challenge: the **spread of fake news**[68,69]. While generative AI has made the production of misleading content easier[70], human-produced misinformation remains highly effective. As a result, the mechanisms described in the previous section (aimed primarily at limiting automated content) are not sufficient to address the broader problem of misinformation.

Vidi must therefore adopt complementary strategies designed specifically to detect, contextualize, and limit the spread of fake news.

A. Collaboration with established media outlets

A first approach is to collaborate closely with established media organizations. Some platforms previously experimented with such partnerships[71,72], notably through **fact-checking** collaborations with professional journalists. These initiatives demonstrated that media organizations can play an important role in identifying and contextualizing misleading information.

However, these collaborations have recently weakened on certain platforms as policy priorities shifted[73], leaving many online spaces less moderated and more vulnerable to misinformation.

Vidi will learn from these earlier initiatives and develop strong, structured partnerships with media organizations. This approach is particularly well suited to the platform's governance model, as public broadcasters (who are central to Vidi's infrastructure) already employ professional journalists and editorial teams with expertise in verification and fact-checking. These existing

---

[68] Aïmeur, Esma, Sabrine Amri, and Gilles Brassard. "Fake news, disinformation and misinformation in social media: a review." *Social Network Analysis and Mining* 13.1 (2023): 30.
[69] Allcott, Hunt, and Matthew Gentzkow. "Social media and fake news in the 2016 election." *Journal of economic perspectives* 31.2 (2017): 211-236.
[70] Kumar, Sanjeev, et al. "Peeping into the future: understanding and combating generative AI-based fake news." *Cognitive Computation* 17.3 (2025): 103.
[71] https://www.lemonde.fr/actualite-medias/article/2017/02/06/huit-medias-francais-s-allient-a-facebook-dans-sa-lutte-contre-les-fake-news_5075054_3236.html
[72] https://partnerships.nature.com/resources/blog/how-to-combat-misinformation-through-collaboration/
[73] https://www.bbc.com/news/articles/cly74mpy8klo

capacities can play a key role in identifying misleading content and providing reliable contextual information.

B. Integrating recent research

Beyond media organizations, collaboration with the academic community could significantly strengthen the fight against misinformation. Over the past decade, researchers[74,75,76,77,78] have developed **sophisticated methods** to study and **detect fake news**.

By maintaining an open platform architecture and providing appropriate research access, Vidi can support these efforts and benefit directly from their findings. Insights from academic research can help improve detection mechanisms, better understand the spread of misinformation, and design interventions that reduce its impact.

Big Tech companies already maintain extensive internal research departments aimed at optimizing their platforms. As a publicly governed social media infrastructure, Vidi should similarly cultivate strong relationships with public research institutions. This collaboration could take the form of joint research programs, open datasets for academic study, or co-funded projects dedicated to improving the safety and resilience of the platform.

C. Contextualization rather than censorship

In addressing misinformation, Vidi will adopt an approach **centered on education** rather than simple censorship. A naive response to fake news might be to remove misleading content entirely. However, a more constructive strategy is to contextualize it and help users understand why it is misleading.

Traditional media frequently apply this approach through fact-checking segments[79,80] that analyze and debunk inaccurate claims. Vidi can follow a similar logic by labeling content identified as false or misleading. These labels would inform viewers while preserving the opportunity to explain and deconstruct the misinformation. For example, recent research have shown that "fact-checking label on misinformation [makes] vaccine attitude more positive."[81] Such scientific results support the generalization of fact-checking labels online.

---

[74] Shu, Kai, et al. "Fake news detection on social media: A data mining perspective." *ACM SIGKDD explorations newsletter* 19.1 (2017): 22-36.
[75] Shu, Kai, Suhang Wang, and Huan Liu. "Beyond news contents: The role of social context for fake news detection." *Proceedings of the twelfth ACM international conference on web search and data mining*. 2019.
[76] Zhou, Xinyi, et al. "Fake news: Fundamental theories, detection strategies and challenges." *Proceedings of the twelfth ACM international conference on web search and data mining*. 2019.
[77] Zafarani, Reza, et al. "Fake news research: Theories, detection strategies, and open problems." *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2019.
[78] Sallami, Dorsaf, and Esma Aïmeur. "Exploring beyond detection: a review on fake news prevention and mitigation techniques." *Journal of Computational Social Science* 8.1 (2025): 23.
[79] https://www.franceinfo.fr/vrai-ou-fake/
[80] https://www.dw.com/en/fact-check/t-56584214
[81] Zhang, Jingwen, et al. "Effects of fact-checking social media vaccine misinformation on attitudes toward vaccines." *Preventive medicine* 145 (2021): 106408.

At the same time, these labels can also be integrated into recommendation systems. Feed generators may choose to reduce the visibility of labeled content, limiting its spread without completely removing it from the platform.

For particularly viral misinformation, Vidi could also notify users who previously interacted with the content to inform them that it has been identified as misleading. Such notifications should be used carefully to avoid overwhelming users and should be designed in collaboration with journalists and editorial experts.

D. Preventing impersonation

Impersonation of trusted institutions, public figures, or well-known organizations have become a common method for spreading misinformation online[82,83,84]. Preventing such impersonation is therefore an essential component of a robust information ecosystem.

Vidi will provide strong verification and authenticity mechanisms allowing institutions, media organizations, brands, and individuals to demonstrate the legitimacy of their accounts. In developing these tools, the platform should draw on recent **scientific advances in digital trust**, such as the Authentimark research initiative[85] currently being developed in the Netherlands.

Addressing impersonation requires a **multi-layered strategy** combining identity verification, account authenticity signals, and technical detection tools. Vidi will continuously integrate state-of-the-art solutions in this domain and treat the prevention of impersonation as a core operational priority.

E. Trusting the community

Finally, as with the fight against AI slop, Vidi will **empower its community** to contribute to the detection and contextualization of misinformation. Community-driven moderation mechanisms have already demonstrated their potential in several online environments.

For example, X's community notes have shown promising results in identifying misleading content. More broadly, large-scale collaborative projects such as Wikipedia demonstrate that online communities are capable of building robust self-moderation systems.

---

[82] Thayer, John. "Defamation or Impersonation? Working towards a Legislative Remdedy for Deepfake Election Misinformation." *Wm. & Mary L. Rev.* 66 (2024): 251.
[83] https://www.newsweek.com/elizabeth-warren-msnbc-republicans-vote-deep-fake-video-1784117
[84] Goga, Oana, Giridhari Venkatadri, and Krishna P. Gummadi. "The doppelgänger bot attack: Exploring identity impersonation in online social networks." *Proceedings of the 2015 internet measurement conference*. 2015.
[85] https://www.uva.nl/en/content/news/news/2025/12/uva-receives-major-nwo-grant-for-groundbreaking-research-to-combat-online-impersonation.html

However, these notes should **only be complementary to other methods**. Research has shown that X's community notes alone are not efficient to combat misinformation[86]. Indeed, the consensus takes time which lets the fake news spread.

Vidi will therefore provide tools that allow users to flag, discuss, and contextualize questionable information. Encouraging these participatory dynamics can strengthen the resilience of the platform while fostering a culture of collective responsibility within the digital public space.

[86] Chuai, Yuwei, et al. "Did the roll-out of community notes reduce engagement with misinformation on X/Twitter?." *Proceedings of the ACM on human-computer interaction* 8.CSCW2 (2024): 1-52.

Chapter IX

User migration: driving adoption through exclusive content

Even the most carefully designed platform can fail if it does not succeed in attracting users. A decisive factor in large-scale adoption is the presence of compelling, distinctive content that users cannot easily find elsewhere. Historically, the success of major platforms has often been triggered by unforeseen events or viral phenomena, such as the role played by Twitter during the Arab Spring. Such **cultural moments**[87] are essential to the success of a social media platform.

For Vidi, however, success cannot depend on chance. Adoption must be deliberately engineered through clear and credible **incentives** for both users and creators to migrate. The primary lever will be the establishment of mechanisms that foster **content exclusivity** on Vidi. As the platform becomes associated with authentic, high-quality, and exclusive content, user migration will follow naturally. This section outlines the complementary strategies designed to drive that transition.

A. Leveraging the catalog of public broadcasters

As Vidi will be operated by public broadcasters, their existing content catalogs constitute a natural foundation for the platform. The **success of broadcaster-owned channels** on platforms such as YouTube, including Arte and similar initiatives, demonstrates their ability to produce high-quality content that appeals to younger audiences.

From the earliest stages, these catalogs will ensure that Vidi offers a rich and attractive viewing experience. Public broadcaster content will provide immediate value to early adopters and guarantee that meaningful discussions and interactions are possible from the first day of operation.

B. Enabling a progressive transition for creators

Vidi is not intended to function solely as a broadcaster-owned streaming platform. Long-term success requires the rapid integration of independent content creators to **diversify the catalog** and stimulate community growth. This transition presents structural challenges: most creators currently depend on existing platforms for both audience reach and revenue and cannot afford an abrupt migration.

Vidi must therefore support a gradual and **low-friction transition.** In a first phase, creators should be able to maintain a **dual presence** across platforms. Tools should be provided to simplify

[87] https://www.newyorker.com/culture/infinite-scroll/social-media-is-navigating-its-sectarian-phase

cross-posting, audience migration, and community management, minimizing operational complexity and reducing the cost of experimentation.

At the same time, Vidi must offer an attractive and **transparent revenue-sharing** model. Higher subscription revenue shares, predictable income mechanisms, and clearer monetization rules can encourage creators to actively redirect their audiences toward the new platform.

Public broadcasters can further accelerate adoption by leveraging existing public subsidies. Content receiving public funding could benefit from temporary exclusivity on Vidi before publication on other platforms. This approach, already used by some subscription-based creators on Youtube, would provide a powerful incentive for both creators and audiences to engage with the platform.

C. European digital rights and freedom at heart: EU Laws effectively enforced

Finally, user migration will be reinforced by the promise of digital rights and freedoms preserved, through regulatory enforcement. Until now, European authorities have been constrained by the absence of credible alternatives. As long as dominant platforms remain indispensable, strict enforcement risks depriving citizens of access to digital public life.

Once Vidi reaches sufficient maturity, this constraint disappears. European states will be able to **enforce digital regulation** more rigorously, without fear of eliminating the only viable platforms. Increased regulatory pressure on non-compliant platforms will affect their economic sustainability and, indirectly, their user experience, accelerating the natural migration toward a healthier and sovereign European alternative.

Chapter X

Community building: designing the Internet's living room

So far, this document has outlined many reasons why users should adopt Vidi. However, rational arguments alone rarely drive the success of online platforms. The internet is an emotional and intuitive environment, where people gravitate toward spaces that feel engaging, welcoming, and socially meaningful. Vidi will not only be the platform people should use; Vidi will be **the platform they love to use.**

To achieve this, Vidi returns to one of the original promises of the internet: community building. Early online spaces were designed to **connect people** across the globe and enable shared experiences. Over time, Big Tech platforms have shifted toward highly individualized and isolating models optimized for engagement metrics. This dynamic is particularly visible on video platforms, where the ideal user is someone endlessly scrolling alone through an infinite feed.

Vidi proposes a different vision. The platform will be designed as a collective space: **the "Internet's living room"** where people can gather, share, and experience content **together**. Watching and exploring online content should once again become a social activity rather than a solitary one.

This collective vision will evolve through many present and future features. In this document, we highlight three core examples that illustrate how Vidi can foster shared digital experiences.

A. Watch parties, creating digital living rooms

Many internet users have experienced the same situation: they want to watch a video, a documentary, or a movie with friends or family members who are not physically present. To do so, they often resort to improvised solutions such as manually synchronizing playback over the phone or relying on unofficial community-built tools.

While these workarounds can be charming in their improvisation, they also reveal an important gap. Despite **clear user demand**, major video platforms have rarely prioritized built-in features that enable shared viewing experiences.

One plausible explanation lies in the dominant attention-based business models of these platforms. Collective viewing tends to reduce the intensity of individual engagement: people talk, pause the video, and focus on their friends rather than continuously consuming content and advertisements. As a group, we don't get isolated and hooked on an infinite feed. As a group, the discussion will "slow down" the watching experience, decreasing the number of ads to be watched. As a group, our attention is available to our watch buddies, not to their ads. From the perspective of advertising-driven platforms, such social dynamics may appear less profitable.

Because Vidi's business model does not rely on maximizing individual attention, it can embrace these social experiences instead of avoiding them. The platform will therefore support integrated “watch parties,” allowing users to watch videos simultaneously, interact in real time, and recreate the feeling of sitting together in the same living room... even when physically apart!

By enabling these shared moments, Vidi aims to transform online viewing into a collective and enjoyable social experience.

B. Embracing reaction live streams

A recent trend on streaming platforms illustrates the growing demand for shared online experiences: the rise of "Just Chatting" content. While Twitch was originally designed around gaming streams, a large portion of the platform is now dedicated to creators interacting directly with their audiences in real time.

Within this category, reaction live streams have become particularly popular. In these sessions, streamers watch and react to videos together with their audience. The format effectively invites viewers into the streamer’s digital living room, creating a collective viewing experience that audiences strongly appreciate.

These shared experiences often lead to meaningful discussions and encourage audiences to engage with content they might not have discovered on their own. Beyond entertainment, they illustrate one of the Internet’s most valuable functions: connecting people around shared cultural moments.

Despite their popularity, reaction streams currently operate in a **legal grey area**. Streamers frequently react to copyrighted material without formal permission from the rights holders. As a result, creators often limit themselves to less prominent content and hope that copyright owners will not intervene. This situation leaves streamers in a precarious position, as copyright claims or account suspensions can abruptly disrupt their activity.

Vidi can provide a more **sustainable framework for these formats**. Because the platform will be supported by public broadcasters and structured partnerships with rights holders, it can help establish fair licensing and revenue-sharing mechanisms for reaction streams. Similar systems already exist in other media industries; for example, the licensing models that allow radio stations, television channels, or public venues to broadcast copyrighted music. Reaction streaming represents a comparable challenge, but one that can be addressed through appropriate legal and economic structures.

C. A community-centered platform

Finally, Vidi should not become a uniform and standardized digital environment. Instead, users should be able to shape and personalize their own digital "living rooms," building communities around shared interests and values.

Existing video platforms present users with a vast and largely homogeneous catalog of content navigated primarily through algorithmic feeds. While efficient for content discovery, this model often discourages the formation of genuine communities.

To encourage richer social dynamics, Vidi will draw inspiration from platforms such as Reddit and traditional online forums. Users will be able to **create and manage their own community spaces**, where videos from selected creators can be gathered and discussed within a shared context.

Over time, online video creation has become increasingly individualized, with creators competing against one another for visibility within algorithmic systems. Yet this was not always the case. In the early 2010s, many creators experimented with collaborative "web TV" projects, pooling their audiences and producing content together. Although many of these initiatives eventually faded, the collaborative spirit behind them remains appealing.

Vidi can help revive this model. By enabling creators to organize around shared communities, the platform can transform isolated channels into vibrant collective spaces where creators support one another and audiences participate more actively in the cultural life of the platform.

Roadmap

Given the exploratory nature of the project, it is not possible at this stage to define a fixed implementation timeline. The pace of development will depend primarily on the level of political commitment and financial support provided by participating European states. Nevertheless, a high-level roadmap can be outlined to identify the main phases required to launch Vidi and progressively scale it into a fully operational European platform.

Phase 1: Coalition formation

Establish a coalition of European countries willing to initiate and politically sponsor the project, with the early involvement of their public broadcasters and relevant authorities. Vidi must be designed from the beginning as a public initiative; it cannot begin as a privately owned company, however well-intentioned. This makes coalition-building a critical step for the platform's long-term success

Phase 2: Budget allocation and legal structuring

Secure a multi-year budget, ideally over a five-year horizon, to guarantee financial stability during the development and launch phases. Create an appropriate legal structure to host the project, drawing inspiration from existing public innovation mechanisms such as the French "start-up d'État" model.

Phase 3: Core team formation

Recruit a core technical and governance team, with a strong emphasis on open-source developers and engineers experienced in decentralized social media ecosystems such as Mastodon, PeerTube, or the AT Protocol.

Phase 4: Minimum Viable Platform

Develop a first operational version of the platform by leveraging existing open-source components whenever possible. This phase focuses on building the core infrastructure, interoperability layers, moderation pipelines, and basic user experience.

Phase 5: Alpha: public broadcaster content only

Open user account creation while initially limiting hosted content to videos produced by participating public broadcasters. This controlled environment will allow the consortium to refine moderation workflows, load balancing, federation mechanisms, and initial recommendation models under real-world conditions.

Phase 6: Beta: verified creators and monetization

Extend hosting to verified independent content creators in order to diversify the catalog and test creator-oriented features. Deploy monetization mechanisms, including subscriptions and, where appropriate, privacy-preserving advertising. Finalize revenue-sharing schemes, strengthen child-protection features, and further refine moderation and recommendation systems.

Phase 7: Official release and public opening

Open the platform to all users, potentially with initial rate limits to manage scale and risk. Expand moderation capacity, optimize infrastructure costs, and stabilize governance and operational processes to support long-term growth and sustainability.

How to support Vidi

If you are a policymaker, a public broadcaster, or anyone interested in supporting this project, please contact Marc Damie or Thomas Yaqoubi Reboul.

Authors

The two authors of this report are technology experts convinced of the importance of building a public service Internet. Having grown up with the Internet, and more specifically with social media. They have experienced both its value and its risks as a digital public space.

They observe that much of today's Internet has been shaped by a generation of founders who built highly influential platforms while benefiting financially from models that have also produced significant societal externalities, particularly for GenZ.

Marc and Thomas believe it is time to rethink and modernize the Internet, giving newer generations a greater role in shaping the digital environments in which they live. *Vidi* reflects an opinionated vision of what a safer, more transparent, and still engaging Internet could look like.

Dr. Marc Damie

Marc is a researcher specialized in privacy-enhancing technologies. His work focuses on improving online privacy, ranging from the identification of privacy vulnerabilities to their mitigation through both technical and socio-technical approaches. He holds a Ph.D. from the University of Lille (France). He conducted his PhD work at Inria Lille (France) and the University of Twente (The Netherlands). Before that, he obtained a MSc. degree in Data Mining from the University of Technology of Compiègne (France). Beyond academia, he has been involved in organizations defending digital rights, notably ctrl+alt+reclaim, a pan-European youth movement he co-founded.

Thomas Yaqoubi-Reboul

Specialized in navigating complex digital public policies at national, European, and international levels, Thomas integrates computer science, data governance, and political science to shape tech landscapes. He honed his statecraft at the Quai d'Orsay, spending three years advising the French Ambassador for Digital Affairs. Driven by a strong commitment to digital rights, he also co-founded ctrl+alt+reclaim with Marc, a civic movement advocating for a better technological future.

Acknowledgments: We want to thank Henri Verdier and Nico Schmidt for their valuable feedback that significantly strengthened the proposal.

Website: https://vidi.damie.eu