Mera News: The Safe Intelligent News Layer for Modern Life

A Privacy-Preserving, AI-Powered Personal News Intelligence System

Mera News B.V., Netherlands

Version 1.0, February 2026

Mera News B.V.
Netherlands
https://mera.news

Abstract

Information overload is the defining challenge of the modern knowledge worker. Professionals spend 60–90 minutes daily filtering noise to find signal, and missing important news carries real financial, career, and personal consequences. Existing solutions force a false trade-off: surrender your privacy for personalization, or stay in the dark.

Mera News eliminates this trade-off. It is a **notification-first personal news assistant** that delivers 3–10 alerts per day, only news that personally impacts you, with clear reasoning and direct links to original sources. All personalization happens on-device, or in Confidential VMs for low-end devices, through a novel privacy-preserving protocol. The architecture mathematically guarantees that no one beyond your own device and its attested Confidential VM can ever access your data.

This paper presents the problem landscape, Mera's architectural innovations, the privacy guarantees that underpin the system, and the value proposition for both readers and publishers.

Contents

1 The Problem: Information Overload Meets Surveillance

1.1 The Scale of the Problem

The average professional today faces an unprecedented volume of information. News is published continuously across thousands of sources, in dozens of languages, spanning topics from local municipal decisions to global economic shifts. The consequences of missing relevant news are tangible:

- A tax policy change that affects your financial planning
- An immigration rule update that impacts your visa status
- A natural disaster in the city where your parents live
- A layoff announcement at your employer
- A school safety alert in your children's district

Yet finding these needles in the haystack requires hours of daily scrolling, scanning, and filtering: time that professionals cannot afford.

1.2 The Privacy Paradox

To personalize news effectively, AI needs to understand deeply personal context: your financial situation, career circumstances, health concerns, family relationships, political views, and geographic connections. This is precisely the information that should never be surrendered to a company's servers.

Every existing AI news product requires exactly that surrender. Users face a binary choice:

- **Opt in to surveillance** → Receive personalised content at the cost of privacy
- **Protect privacy** → Accept generic, irrelevant content and spend hours filtering manually

Mera proves this is a false dichotomy.

1.3 The Broader Context

This problem exists within a larger crisis of trust in information systems:

- **Trust in media is at historic lows.** Users increasingly recognize they are the product, not the customer.
- **Information manipulation is deliberate.** Noise is often manufactured to drown out signal by governments, corporations, and bad actors.
- **Engagement-maximizing algorithms cause harm.** Average screen time now exceeds sleep time. The backlash against attention-extracting design is real and growing.
- **AI summary layers threaten publishers.** Readers get the gist without ever seeing the source, collapsing the economics of quality journalism.

2 Existing Solutions and Their Failures

Category	Examples	Core Failure
News aggregators	Apple News, Google News, Flipboard	Feed-based engagement farming. Generic. Optimized for time-on-app, not value delivered.
AI news tools	Artifact (shut down), Syft AI, ChatGPT Pulse	Cloud-centric. Privacy is a policy, not an architecture. Summary-first disintermediates publishers.
Keyword alerts	Google Alerts	Literal keyword matching. No semantic understanding. High noise, duplicates, missed relevance.
Social media	Twitter/X, Reddit, LinkedIn	Algorithmic manipulation. Doomscroll design. Surveillance as business model.
Premium subscriptions	NYT, WSJ, Bloomberg	Single-source. Expensive. Still generic: same content for every subscriber.
RSS readers	Feedly, Inoreader	Manual curation burden. No AI filtering. No deduplication. No proactive alerts.
Privacy-focused	Brave News, Kagi News	Better privacy, but feed-based (requires browsing). No proactive notifications. No personal context.

No existing solution combines: proactive notification delivery, deep personal context understanding, privacy-by-architecture, source attribution, and explicit per-alert reasoning.

3 The Mera News Solution

3.1 Core Product

Mera News is a **notification-first personal news assistant.** The product design is deliberately minimal:

1. **You tell Mera what matters:** your job, investments, family, interests, concerns, in natural language.
2. **This stays on your device**, encrypted, never uploaded.
3. **Mera scans hundreds of sources hourly**, using AI to filter thousands of articles.
4. **You get alerts at your configured times every day**, each with a "Why this matters to you" explanation.
5. **One tap → original source**: we route to publishers, not replace them.

3.2 Key Differentiators

Notification-first, not feed-first. Mera succeeds when you spend *less* time in the app. There is no feed to scroll, no engagement loop to exploit. You receive precisely the news that matters, then you're free to live your life.

Explicit reasoning. Every alert includes a natural-language explanation of why this specific article matters to your specific situation. Users can evaluate and provide feedback, creating a learning loop.

Source-agnostic and publisher-friendly. Mera displays headlines only: no summaries, no scraped content. Every tap sends users directly to publisher websites. Mera is a traffic driver, not a content aggregator.

Privacy that's provable, not promised. Mera's privacy guarantees are architectural, not contractual. The system is designed so that even Mera's own employees cannot access user data. The client code is open source, and anyone can verify these claims.

3.3 How Personalisation Works

Mera's personalisation operates through a three-stage intelligence pipeline:

Stage 1: Fact Extraction & Topic Generation. An on-device LLM first extracts personal **facts** from the user's free-text profile (e.g., "user lives in Amsterdam", "user's parents live in Kathmandu"). These facts never leave the device. The LLM then uses these facts to generate 20–30 location-anchored **topics**, precise news queries like "Amsterdam cycling infrastructure" or "Nepal earthquake alerts", organized by priority tiers (safety, direct impact, relevant interests, informational, awareness). Only topics (with noise added) are sent to the server.

Stage 2: Multi-Source Article Matching. The server continuously ingests articles from hundreds of RSS feeds across languages. Each article is language-detected, translated to English, embedded into a unified vector space using `nomic-embed-text-v1.5`, and clustered with related articles to form coherent news stories. This cross-language clustering means a Hindi article and an English article about the same event are grouped together.

Stage 3: Contextual Relevance Scoring. The user's interest embeddings are matched against article and cluster embeddings. Candidates are then scored by an LLM

against the user's full profile on a 0.0–1.1 scale, where 1.1 is reserved for immediate life-threatening emergencies that override all notification preferences. Each scored article receives a natural-language reason explaining its relevance.

4 The Mera Protocol: Privacy-Preserving Personalisation

4.1 Design Philosophy

The Mera Protocol is a privacy-preserving hybrid architecture for client-server AI applications. It provides mathematical privacy guarantees while enabling efficient server-side processing.

Core insight: The rich user profile and extracted facts stay local. The server sees only noisy, k-anonymous topics: broad news queries indistinguishable from thousands of other users.

4.2 Protocol Overview

The protocol operates in three phases:

Phase 1: Fact Extraction & Topic Generation (Client-Side). After the user creates their profile, an on-device LLM first extracts personal **facts** (e.g., "user lives in Eindhoven", "user's parents live in Kathmandu") that stay on-device. It then generates 20–30 **topics** from these facts: broad news queries that get noise-padded before being sent to the server:

```
// Facts (NEVER leave device)
const facts = [
"User lives in Eindhoven, Netherlands",
"User is a Nepal citizen, software engineer at Meta",
"User is on 30% ruling for 2 more years",
"User's parents are old and live in Kathmandu",
"User is dating a German woman"
]

// Topics generated from facts (noise added before sending)
const topics = [
"Eindhoven local news and events",
"Netherlands immigration and visa policy changes",
"Meta company news and announcements",
"Kathmandu safety and emergency alerts",
"Netherlands-Germany cross-border news",
// ... 20-30 total
]
```

Privacy noise is then added, both fake locations and expanded categories, to achieve k-anonymity:

```
// Noisy topics (sent to server)
const noisyTopics = [
// Real topics
"Eindhoven local news and events",
"Netherlands immigration and visa policy changes",
"Meta company news and announcements",
"Kathmandu safety and emergency alerts",
// Noise topics (indistinguishable from real)
```

```
"Amsterdam startup ecosystem news",
"Berlin tech industry developments",
"Rotterdam port and shipping industry",
// ... server cannot tell which are real vs noise
]
```

Phase 2: Retrieval (Server-Side, Continuous). The Mera backend fetches and preprocesses news articles from hundreds of sources, matching them against user topic sets. The volume of matched articles varies based on topic coverage and news activity.

Phase 3: Filtering and Scoring (Client-Side). Every hour, the device pulls matched articles (volume varies based on topic coverage and news activity). The on-device LLM then:

1. Strips noise topics: keeps only articles matching original topics
2. Detects novelty: avoids repeating known information
3. Scores with full facts context: uses extracted facts to understand *why* each article matters
4. Generates reasons: creates a natural-language explanation for each alert
5. Selects top 3–10: delivers as notifications

4.3 Privacy Guarantees

The protocol provides strong privacy properties:

Property	**Guarantee**
Profile & fact confidentiality	User profile and extracted facts never leave the device. Server receives only noisy topics: broad news queries with no personal context, employment, family, or financial details.
Location ambiguity	Server cannot distinguish real locations from noise. "Does the user live in Eindhoven or Amsterdam?" is unanswerable.
Category ambiguity	Server cannot determine which categories are genuine interests vs. noise padding.
Notification opacity	Server sends matched candidates. Which 3–10 become notifications is decided on-device. Server never learns what the user actually saw.
k-anonymity	Each noisy topic set matches thousands of other users. No individual can be singled out from their query pattern.
Ephemeral processing	No user data is written to server databases. Interest maps are processed in-memory and immediately discarded.

4.4 On-Device LLM

The Mera Protocol requires an on-device LLM to handle privacy-preserving query generation and response personalization. The system uses a **LoRA fine-tuned Llama 3.2**

3B model (~1.5 GB base + ~50 MB adapter), optimized for two core actions:

- **Extract & Anonymize:** Extract facts from profile, generate topics from facts, and add noise topics for k-anonymity.
- **Personalize:** Filter batch results using extracted facts locally, scoring relevance and generating explanations.

The adapter is trained on ~2,000+ examples covering diverse profiles, trait distribution, and edge cases.

5 Technical Architecture

5.1 System Overview

Mera is built as a NestJS monorepo comprising three microservices communicating through a shared MongoDB database and Redis-backed job queues:

- **GraphQL API**: User-facing service for personalized suggestions, profile management, and job triggers.
- **Auth Service**: Passwordless email OTP authentication.
- **Async Engine**: Background job processor handling 28 BullMQ queues for news ingestion, AI processing, personalization, and notifications.

5.2 News Processing Pipeline

News flows through a five-stage pipeline:

1. **Ingest**: Fetch articles from multilingual RSS feeds.
2. **Detect Language**: Identify actual content language (RSS metadata is unreliable).
3. **Translate**: Convert non-English articles to English, preserving originals.
4. **Embed**: Generate vector embeddings from English text using `nomic-embed-text-v1.5` (768 dimensions). Title + summary (50–70 words) produces optimal topic similarity.
5. **Cluster**: Group similar articles into coherent stories via vector similarity (threshold ≥ 0.9). Articles auto-expire after 24 hours.

The result: a continuously updated, multilingual news corpus organized into coherent story clusters, ready for personalization.

5.3 Key Technical Innovations

Interests as semantic queries, not categories. "Amsterdam cycling infrastructure" matches news about cycling in Amsterdam but not Tokyo. This location-anchoring is fundamental to Mera's relevance.

Hierarchical location awareness. Mera models the user's relationship to geography as a multi-layered graph: home city, partner's family city, country of origin, workplace, each with relationship-type weights and distance-aware scoring.

Multi-language unified embedding. Translating to English before embedding creates a single vector space where stories in different languages cluster naturally. A Hindi report and an English report about the same event produce similar vectors.

Life-safety override (Score 1.1). The relevance scale intentionally exceeds the normal 0.0–1.0 range. A score of 1.1 ensures immediate life-threatening emergencies always surface above all other content, regardless of notification preferences.

Article-centric bidirectional matching. Traditional systems match clusters to interests in one direction. Mera also matches individual articles to interests, then uses those matches to *promote* their parent clusters, surfacing relevant content that pure cluster-level matching would miss.

6 The Confidential VM Fallback Path

Not all devices can run a 3B-parameter LLM locally. Mera provides a **Confidential VM fallback** for low-end devices that preserves privacy guarantees through hardware-level isolation.

6.1 Design Principles

- **Protocol symmetry.** The Mera backend sees identical traffic from both on-device and Confidential VM users. The two paths are indistinguishable server-side.
- **Zero trust coupling.** The Confidential VM never interacts with the Mera backend. It only communicates with the user's device.
- **Stateless compute.** The VM spins up, decrypts, computes, returns results, wipes memory.

6.2 Two-Key Encryption

The user's profile is encrypted using a **threshold decryption** scheme:

- **User key:** Derived from a PIN using Argon2id (never stored).
- **System key:** Held by Mera, released only to attested Confidential VMs.
- **Neither party alone can decrypt.** The combined encryption key exists only inside the TEE during computation.

6.3 Data Flow

Setup (once + rare profile updates): Because the device cannot run the LLM locally, it sends the encrypted profile and user key to a Google Confidential Space VM (AMD SEV-SNP) for processing. The VM decrypts the profile, extracts facts from the profile, generates topics from facts, adds noise to topics, and returns the noisy topics to the device. The device then sends the noisy topics to the Mera backend, identical to the on-device path. The VM wipes all decrypted data.

Continuous ingestion (server-side, no user interaction): The Mera backend fetches and preprocesses news articles, matching them against stored user topic sets.

Scoring (every hour): The device pulls matched articles from the backend (volume varies based on topic coverage and news activity). Because the device cannot run the LLM locally, it sends the encrypted profile, articles, and user key to the Confidential VM for processing. The VM decrypts, scores, ranks, generates reasons, and returns results. The device displays the top 3–10 articles. The VM wipes all data.

6.4 Privacy Comparison

Guarantee	On-Device	Confidential VM
Profile never visible to Mera	✓ Mathematical	✓ TEE isolation
Profile never visible to third party	✓	✓ Hardware attestation
Backend sees identical traffic	✓	✓ Protocol symmetry
Trust model	Zero trust	Trust in hardware TEE
Offline capable	✗ Requires network	✗ Requires network

7 Publisher Value Proposition

Mera is designed as a **distribution partner that keeps publishers as the destination** in an AI-mediated world.

7.1 Core Alignment

- **Pure referral traffic.** Users click through to publisher sites, generating full pageviews, not snippets or summaries.
- **Monetization stays with publishers.** Ads, paywalls, subscriptions, newsletters. Mera doesn't interfere.
- **No integration required.** Publishers don't need to sign up, build APIs, or change anything.
- **High-intent readers.** Users who click a Mera alert *want* to read that article. This is qualified traffic.

7.2 Discovery Engine

Smaller or niche publishers surface alongside major names based on user interests, not brand recognition. A hyperlocal Dutch outlet covering Eindhoven zoning decisions reaches the user who lives there, something Apple News or Google News would never surface.

7.3 Partnership Programs

Referral attribution. Monthly traffic reports showing Mera-delivered readers, click-through rates by topic, and conversion lift.

Paywall partnership. Mera users get 3–5 free articles/month from partner publishers, triggered only when Mera surfaces an article as *personally relevant*. This converts better because users are pre-qualified: they know the article matters to them.

Co-marketing on privacy. Joint narrative: "Quality journalism deserves quality distribution, without surveillance."

Local news partnerships. Local news is dying, but local readers still exist. Mera surfaces hyperlocal stories that big aggregators ignore.

7.4 Publisher Credibility and Trust Signals

Mera leverages established journalism standards to help users assess source quality:

- **Journalism Trust Initiative (JTI)**: ISO-style certification by Reporters Without Borders
- **The Trust Project**: 8 Trust Indicators adopted by hundreds of outlets globally
- **NewsGuard**: Reliability ratings based on 9 apolitical criteria
- **IFCN**: Poynter's accreditation for fact-checking organizations

Badges are displayed alongside articles. Fact-check pairings show both the news and the fact-check together. The approach keeps Mera open while giving users context to make informed choices.

The complete list of publishers on Mera is open source and publicly available at https://github.com/abhijeet1403/news-feed-list-of-countries.

8 Trust and Credibility

8.1 The Credibility Challenge

Every company claims to be "privacy-first." Users have heard it all before. Mera's answer: make privacy claims **independently verifiable** so users don't need to take our word for it.

8.2 Verification Toolkit

Mera provides multiple verification pathways:

- **Open-source client code** on GitHub from day one. Users can compile from source and verify it matches the App Store version.
- **Network traffic analysis guides**: users can capture Mera's traffic with standard tools and see exactly what's sent to servers.
- **Independent security audits**: pre-launch and biannual, by different firms.
- **Reproducible builds**: cryptographic proof that the distributed binary matches the public source.
- **Annual transparency reports**: government data requests, security incidents, and privacy metrics.
- **Bug bounty program**: rewarding discovery of critical findings.

9 Business Model and Market Opportunity

9.1 Revenue Model

Subscription-only. No ads. No data monetization. Revenue is generated when users find Mera valuable, not when data is extracted.

Tier	Features
Basic	24 notification windows/day, standard personalisation, all sources, limited story tracking
Professional	Expanded alerts, deeper analysis, comprehensive story tracking

9.2 Market Timing

Seven converging trends make this the right moment:

1. **AI and privacy awareness peaked.** Users understand AI can personalize; privacy concerns are mainstream.
2. **Privacy regulations tightening.** GDPR, DMA, DSA favor architectures like Mera's.
3. **On-device AI finally viable.** Apple Neural Engine and Qualcomm NPU make local inference practical.
4. **News subscription market proven.** NYT has 10M+ subscribers; Proton has 100M+ paying for privacy.
5. **Trust in information systems at historic lows.** Users seek alternatives to surveillance-based platforms.
6. **Information overload reaching epidemic levels.** People are exhausted and actively seeking solutions.
7. **Growing awareness of information manipulation.** Noise is recognized as deliberate, driving demand for signal.

9.3 Why Big Tech Won't Build This

Factor	Why Mera Wins
Business model conflict	Big Tech's ad revenue requires data collection. Privacy-first personalization is antithetical to their economics.
Verifiability	Mera's open-source client proves privacy claims. Big Tech operates as a black box.
Focus	News is their side project. It's Mera's entire company.
Model improvements are a tailwind	Better on-device models amplify Mera's core promise: more accuracy, same privacy.

10 Conclusion

Information overload is a solved problem, if you're willing to surrender your privacy, or spend significant resources building custom solutions on today's generic LLMs.

Mera News proves you don't have to choose.

The Mera Protocol demonstrates that deep personalization and rigorous privacy are not competing goals; they are complementary when the architecture is designed correctly. On-device AI processes the user's full context locally. The server sees only k-anonymous queries indistinguishable from thousands of others. For devices that can't run local inference, a Confidential VM provides an architecturally equivalent fallback. In both paths, the user's profile is never visible to Mera or any third party.

For publishers, Mera is the distribution partner the industry needs: one that routes qualified readers to original sources, respects paywalls, and makes niche reporting economically viable through intent-matched delivery.

The safe intelligent news layer for modern life.

For inquiries, contact contact@mera.news.

Mera News B.V.
Netherlands
https://mera.news