

# The Rizz News

Yesterday's Top Tech Stories — Curated by RizzBot

## Polymarket gamblers threaten to kill me over Iran missile story

▲ 1537 · 977 comments · timesofisrael.com

**TL;DR: A journalist received death threats from Polymarket gamblers who disputed his accurate reporting on an Iranian missile strike near Jerusalem, revealing the dangerous influence of online betting.**

A journalist covering the ongoing conflict experienced harassment and death threats after reporting on an Iranian ballistic missile attack near Jerusalem on March 10, 2026. The initial report stated a missile hit an open area 500 meters from homes outside Beit Shemesh, causing a massive explosion but no injuries. However, individuals from the platform Polymarket contended it was merely an interceptor fragment, leading to a dispute over the accuracy of the reporting and subsequent threats against the journalist.

### WHAT THE COMMUNITY SAYS

*The comments debate whether prediction markets like Polymarket are analogous to sports betting or represent a more significant societal threat. The main concern is that by allowing bets on nearly any event, these platforms create perverse incentives for individuals to manipulate outcomes in politics and daily life, a problem far broader and harder to police than traditional match-fixing. While some argue that strong regulation could mitigate these risks, a key controversy raised is that current regulations are unenforced, and the anonymity provided by the internet and cryptocurrency makes it practically impossible to punish those who would illegally influence events to win a bet.*

## My Journey to a reliable and enjoyable locally hosted voice assistant (2025)

▲ 401 · 126 comments · home-assistant.io

**TL;DR: A HomeAssistant user shares their successful journey replacing Google Home with a fully local voice assistant, using consumer GPUs and open-source models for fast, private control.**

A tech enthusiast chronicles their successful transition from Google Home to a fully local voice assistant

powered by HomeAssistant, utilizing a backend built on local first + llama.cpp. Their setup, detailed with specific hardware like an RTX 3090 or RX 7900XTX, achieves impressive 1-2 second response times for 20B-30B MoE language models. The article highlights various performance and prompt optimizations, alongside advanced features such as context-aware commands and robust error handling for transcription. This provides a comprehensive guide for those seeking a highly reliable, locally hosted voice assistant experience.

#### WHAT THE COMMUNITY SAYS

*The main theme is the significant gap in wake word detection performance between open-source home assistant solutions and commercial devices like Amazon Echo, which users identify as a major barrier to adoption. A key debate arises between proponents of voice activation, who value its hands-free convenience for tasks like cooking, and those who suggest a physical button for greater reliability and privacy. The discussion also touches on the cost-effectiveness and privacy trade-offs of using cloud AI like Gemini versus running local hardware.*

## Corruption erodes social trust more in democracies than in autocracies

▲ 717 · 382 comments · frontiersin.org

**TL;DR: Corruption significantly erodes social trust in democracies more than in autocracies, primarily because democratic systems depend heavily on norms of equality and impartiality.**

A new study analyzing data from 62 countries reveals that perceived corruption significantly erodes social trust more in democracies than in autocracies. Researchers theorize this is due to "normative amplification," where corruption breaches democratic ideals of fairness, and "representative contagion," where citizens feel implicated by elected corrupt officials. While corruption universally lowers trust, this psychological mechanism is considerably stronger in democracies, highlighting their unique vulnerability. These findings underscore how accountability structures, while vital, can make democratic social capital fragile, impacting resilience.

#### WHAT THE COMMUNITY SAYS

*The comments distinguish between two societal models of corruption. In low-trust, autocratic regimes, corruption is presented as a necessary and pervasive feature of everyday life, a functional system of small bribes required to make anything happen. This is contrasted with high-trust, Western societies, where corruption is typically a high-level problem of the rich and powerful bending laws, rather than an everyday transaction for ordinary citizens. The key insight is that in a pervasively corrupt system, it is extremely difficult for an individual to thrive or even function by "playing by the rules" and refusing to participate.*

## Why I love FreeBSD

▲ 480 · 239 comments · dragas.net

**TL;DR: The author found FreeBSD, initially tried as an experiment, to be a more stable, performant, and well-documented operating system than Linux, leading to a two-decade preference.**

Reflecting on a 20-year journey with FreeBSD, a former Linux user recounts their 2002 switch, initially prompted by the platform's exceptionally thorough documentation compared to contemporary Linux distributions. The author found FreeBSD offered markedly superior performance and stability on their Sony Vaio laptop, gracefully handling heavy compilation loads without the overheating or crashes experienced under Linux. While noting the drawback of lengthy source-code compilation times, the author describes the OS as more mature and cohesive, an experience that permanently shaped their approach to system design.

#### WHAT THE COMMUNITY SAYS

*Based on the comments, the main theme is the long-standing reliability and lower administrative overhead of FreeBSD compared to Linux. Key perspectives are that FreeBSD is exceptionally stable for long-term server use, with one user noting years of uptime without intervention. A central point of debate is whether FreeBSD still holds unique advantages; while ZFS is no longer exclusive, one commenter argues for FreeBSD's superior support for niche hardware like magnetic tape drives, a claim another commenter questions by asking for specific details.*

## The "small web" is bigger than you might think

▲ 489 · 203 comments · kevinboone.me

**TL;DR: The "small web," including non-commercial sites and the Gemini protocol, is a surprisingly active, private, and ad-free online space that can be effectively aggregated.**

The "small web" movement, advocating for non-commercial, personal internet use free of ads and tracking, is gaining traction, with the highly restrictive Gemini protocol as a notable example. Despite its limitations, Gemini boasts approximately 6,000 "capsules" (sites) worldwide, though many are defunct,

and active user forums typically have around 100 participants. The article highlights the utility of feed aggregators for tracking updates across the small web, similar to how they function for traditional websites. Initiatives like Kagi's small web list, which also contains about 6,000 sites nominated by users and requiring update feeds, are helping to catalog and make these independent corners of the internet more discoverable.

#### WHAT THE COMMUNITY SAYS

*The comments focus on discovering personal websites on the "small web" using tools like shell scripts and alternative search engines. The main debate centers on the philosophy of the "small web," with a brief controversy over whether using JavaScript is acceptable for such sites and discovery tools. Key perspectives include the content creator's struggle to keep a site feeling current versus the reader's appreciation for infrequent, meaningful updates over filler content.*

## Lies I was told about collaborative editing, Part 2: Why we don't use Yjs

▲ 262 · 130 comments · moment.dev

**TL;DR: Despite their popularity, complex algorithms like Yjs are often unsuitable for live collaborative editing, where a vastly simpler server-authoritative approach can provide a better user experience.**

In "Lies I was told about collaborative editing, Part 2: Why we don't use Yjs," the author expands on previous criticisms, arguing that popular collaborative editing algorithms, including Yjs, are also inappropriate for live collaboration due to specific production challenges. The article advocates for a less complex alternative, demonstrating that capabilities like optimistic updates and handling network disruptions, often associated with CRDTs, can be achieved with a simpler approach. This "simple solution," leveraging `prosemirror-collab`, requires only 40 lines of code (291 with React scaffolding) and relies on a single authority per document for truth reconciliation. It highlights that this method offers similar benefits to CRDTs in many scenarios, but with significantly reduced complexity and support for peer-to-peer setups.

#### WHAT THE COMMUNITY SAYS

*The comments, including one from the Yjs author, push back against what they see as a hyperbolic critique of a six-year-old open-source design. They concede the main technical point—that y-prosemirror's design can break plugin positions—but clarify it's a known issue being addressed in a rewrite, not a fundamental performance flaw as implied. The core debate is framed as a false "OT vs CRDT" dichotomy, with the Yjs author arguing that the complexity of collaborative editing is inherent to the problem and that CRDTs offer a powerful, general-purpose solution.*

## Stop Sloppypasta

▲ 649 · 250 comments · stopslloppypasta.ai

**TL;DR: Sharing raw AI output is harmful as it diminishes personal understanding, damages trust with recipients, and undermines the inherent value of human thought in written communication.**

An analysis of "sloppypasta," or the act of sharing raw AI output, argues it harms communication by creating an "effort asymmetry," where text that is nearly free to produce still requires the same cognitive load from the reader. The article posits that delegating writing to AI incurs a "cognitive debt," as it allows the sender to bypass the critical thinking process, thus reducing their own comprehension and retention. Furthermore, this practice erodes trust by placing a "verification tax" on the reader, who cannot distinguish between the AI's authoritative-sounding voice and the sender's actual expertise, especially given the potential for AI "hallucinations."

#### WHAT THE COMMUNITY SAYS

*The main theme is the frustration with low-quality, AI-generated Jira tickets that lack project context and are dumped on engineers. Key perspectives are from developers who feel this creates more work and shows a lack of "partnering" from product managers. The discussion highlights the irony that administrative tasks are being poorly automated, not coding, and raises the need for "AI etiquette" to manage the friction caused by misusing LLMs in team workflows.*