

Forensic Disclosure Prep Brief: Submitting an Amazon Fire TV / SK Hynix eMCP Hardware-Security Report to KrebsOnSecurity

TL;DR

- **Submit via Brian Krebs's stated secure channel** — **Signal username** `briankrebs.07` — **with the public tip email** `krebsonsecurity@gmail.com` as backup; there is no published PGP key and no formal hardware-disclosure portal, so **frame the report as an investigative-journalism tip, not a vendor bug-bounty submission**. Krebs has no security.txt-style intake; he recommends security.txt for *organizations*, but for himself relies on email plus Signal.
- **The strongest, most defensible technical core is the eMCP/RPMB attack literature** (Fukami & Buurke EMFI, WiSec '24; Fukami/Buurke/Geradts key-extraction, ePrint 2024/180; CERT VU#231329) plus the FCC confidentiality regime (47 CFR §0.457/§0.459) and Amazon's documented shell-entity filing practice ("Vlado L.L.C." = FCC grantee 2AE6S = Fire TV Stick LY73PR is confirmed). These are publication-grade.
- Your "poisoned training corpus via FCC documentation gaps" thesis is novel as a *synthesis* but **NOT** as primitives: data-poisoning, long-tail-knowledge hallucination, and FTC AI-substantiation enforcement are all separately well-documented. Treat the FCC-gap-causes-AI-misID claim as an original hypothesis you are advancing, clearly labeled as such, not as established fact.

Key Findings

1. KrebsOnSecurity contact, disclosure policy, and preferred format

- **Primary tip email:** `krebsonsecurity@gmail.com` (published on the "About this Blog" page via a mailto link and a web contact form with captcha). `krebsonsecurity`
↗
- **Secure channel:** Signal username `briankrebs.07`, published verbatim in his verified Mastodon bio at `infosec.exchange/@briankrebs` ("Signal: briankrebs.07 krebsonsecurity @ gmail .com") `Infosec Exchange`
↗ and mirrored on Muck Rack. This is a Signal *username*, not a phone number.
- **No published PGP/GPG key** was found on the About page, the Mastodon bio, or elsewhere; the "About the Author" page offers only a contact form and invites reader correspondence.
- **No formal disclosure policy / no hardware bug-bounty intake.** Krebs is an investigative journalist, not a vendor. In his 2021 security.txt piece he notes self-appointed pen-testers flood tip channels with low-quality automated scan reports, and that such "nuisance reports can be ignored or grouped" with filters. Implication: a hardware-security report must immediately signal that it is original, verifiable, and newsworthy (named entities, primary documents, reproducible method) to clear his spam filter.
- **Preferred format (inferred from his published work):** Krebs's investigations center on named entities, direct documentary evidence, and on-the-record responses from the accused party (he routinely contacts subjects for comment before publishing). A submission mirroring that structure — a tight narrative, primary-source exhibits (FCC filing IDs, datasheets, chip photos), and a clear "why this matters to consumers" — fits his house style.

2. Prior Krebs coverage of Fire TV hardware / FCC opacity / eMCP misID

- **No KrebsOnSecurity article was found specifically on Fire TV hardware security, FCC filing opacity, or eMCP misidentification.** This is a genuine coverage gap — which is favorable for pitching novelty.
- Closest adjacent Krebs work: the 2021 “Does Your Organization Have a Security.txt File?” post (disclosure-channel mechanics) and ongoing IoT/botnet coverage (e.g., the Kimwolf [Muck Rack](#) [IoT botnet](#), 2026).
- The authoritative Fire TV FCC-leak beat lives elsewhere — **AFTVnews** (Elias Saba) and **Liliputing** systematically track Amazon’s shell-company FCC filings. The seminal *security* research on Amazon device memory is **Giese & Noubir, “Amazon Echo Dot or the Reverberating Secrets of IoT Devices,” WiSec ‘21**, which extracted Wi-Fi credentials, tokens, and prior-owner data from used Echo Dots’ eMCP flash even after factory reset (wear-leveling plus lack of encryption); they explicitly worked with **Hynix eMCP** packages and disclosed to Amazon.

3. FCC confidentiality regime (47 CFR §0.457 and §0.459) and AI pipelines

- **§0.457** lists categories of records the Commission withholds by statute (trade secrets, privileged commercial/financial info).
- **§0.459** governs *requests* to withhold materials from public inspection: each request must identify the specific info, the proceeding, the degree to which it is commercial/trade-secret, a competitive-harm analysis, measures taken to prevent disclosure, and whether it is already public. Denials are appealable within 10 business days. [Legal Information Institute](#) [FCC](#)
- **Equipment-authorization specifics:** Under the FCC’s confidentiality-request procedures (KDB 726920 D01), **schematics, block diagrams, operational descriptions, parts lists (BOM), and tune-up procedures** are eligible for **permanent (long-term) confidentiality**; [FCC](#) [Acbscert](#) [internal photos, external photos, test-setup photos, and user manuals](#) can be held under **short-term confidentiality up to 180 days** from grant (commonly used to mask unreleased products). Test reports, RF-exposure reports, external photos, and FCC-ID labels generally cannot be permanently hidden. [FCC](#)
- **A public-notice principle matters for your thesis:** the FCC states that information already public “is not confidential” and won’t be granted confidential treatment. [FCC](#)
- **Interaction with AI training pipelines:** No FCC rule addresses AI training. The interaction is *de facto*, not regulatory: the public EAS/FCC-ID corpus (mirrored by [fccid.io](#), [fcc.report](#), [fcc.io](#)) is heavily scraped and is a plausible training/RAG source for hardware Q&A. When a manufacturer redacts BOM/schematics/internal photos under §0.459, the authoritative chip-identity record for that device is *absent from the public corpus*. This is the documentation gap your thesis builds on. (This causal chain is your original argument; the FCC rules themselves are silent on it.)

4. Security research on SK Hynix H9TQ64A8GTDC eMCP specifically

- **No CVE is assigned to the H9TQ64A8GTDC part specifically.** It is a real SK Hynix eMCP: **8GB eMMC 5.1 + 8Gb LPDDR3, 221-ball FBGA**, [indasina](#) [full part number H9TQ64A8GTDCUR-KUM](#) (confirmed via the ElneC programmer DB, DigiPart, and distributor listings). Note the naming-collision hazard your report should flag: some distributor/aggregator pages mislabel the sibling H9TQ64A8GTAC variant as “a dram chip ... 8gb lpddr3” [Datasheets.com](#)

— i.e., the eMCP (combined NAND+DRAM) gets reduced to a DRAM-only description. That is exactly the class of misidentification your thesis concerns.

- **Relevant device-class CVEs / advisories (not part-specific):**
- **CERT/CC VU#231329** — RPMB protocol does not adequately defend against replay attacks (eMMC/UFS/NVMe broadly); references Western Digital and Intel SA-00391.
- **CVE-2025-6202 (“Phoenix” Rowhammer)** — affects SK Hynix **DDR5** (not eMMC), CVSS 7.1, discovered by ETH Zürich COMSEC with Google. Disclosed via the Swiss NCSC on June 6, 2025, with embargo lifted September 15, 2025. All **15 tested SK Hynix DDR5 DIMMs** produced 2021–2024 were vulnerable; it is the first privilege-escalation exploit “running in default settings on a PC in as little as 109 seconds,” achieved RSA-2048 key extraction on 73% of DIMMs, and the only validated mitigation is a 3× refresh rate at ~8.4% performance cost. To be presented at IEEE S&P 2026. Relevant to “SK Hynix memory security” framing but a different product line from the eMCP.
- **Extraction / RPMB-key research (directly on point, eMMC eMCP):**
- **Aya Fukami, Richard Buurke & Zeno Geradts (2024), “Exploiting RPMB authentication in a closed source TEE implementation”** (Cryptology ePrint Archive Paper 2024/180; published in *Forensic Science International: Digital Investigation*, DOI 10.1016/j.fsidi.2023.301682): the RPMB authentication key can be **extracted directly from the target memory chip**, then used to defeat anti-rollback [ScienceDirect](#)
↗
and restore wiped data.
- **Fukami & Buurke (2024, updated v1 arXiv 2511.22340 submitted Nov 27, 2025), “Keyless Entry: Breaking and Entering eMMC RPMB with EMFI”** (WiSec ‘24, DOI 10.1145/3643833.3656114; authors affiliated with the Netherlands Forensic Institute): electromagnetic fault injection on **three eMMCs from a single major manufacturer (Samsung eMMC controllers)** skips the RPMB HMAC check; RPMB authentication was glitched and data in **two of the three eMMCs** was “overwritten with arbitrary data, without affecting the integrity of other data.” The authors state, “The total cost of the equipment, needed to re-create the setup, is less than USD\$7,000, which makes it a feasible option for most attackers,” and used **publicly known vendor-specific commands** to get code execution on the eMMC controllers. [ResearchGate](#)
↗
Important honesty caveat for your report: this EMFI study’s targets were Samsung eMMC controllers, not SK Hynix — do not imply it was performed on the H9TQ64A8GTDC.
- Western Digital’s RPMB protocol-vulnerabilities white paper documents two replay attacks against the protocol.
- **RPMB background:** 32-byte key, HMAC-SHA256, write-once provisioning, monotonic counter; [Skyhighmemory](#)
↗
meant to be owned by a TEE. [arxiv](#)
↗
If provisioning/secret handling is weak (as in the closed-source TEE studied), forensic key extraction is feasible.

5. Amazon FCC shell-entity filings and criticism

- **Confirmed Amazon FCC proxy: “Vlado L.L.C.” — FCC grantee code 2AE6S — FCC ID 2AE6S-0948, an “HDMI Digital Media Receiver” = Amazon Fire TV Stick (model LY73PR), approved Sept 14, 2017.** ([fccid.io](#) / [fcc.report](#); device identity corroborated by the Kodi wiki, [deviwiki](#), and [AFTVnews](#).)
- **The shell-company practice is well-documented and routine.** [AFTVnews](#): “As with all Amazon hardware, the FCC filing is done using a shell company.” [AFTVnews](#)
↗
Liliputing: Amazon “registers products under a shell company ... using PO boxes and VoIP phone numbers, and

generally creating companies that don't look like real companies." Liliputing

↗

- **Other confirmed Amazon FCC shell names** (corroborating the pattern): Quill Royal LLC (Fire TV, Boise PO box), Liliputing

↗

Bluff Arches Cay LLC (Fire TV remote), AFTVnews

↗

Abyssal Plain LLC (grantee 2AWRO, Fire tablets), Liliputing

↗

Polebridge Hole LLC, Liliputing

↗

Blancopage LLC, Triesan LLC, Liliputing

↗

Jib Wresh LLC, Liliputing

↗

AzulFlower LLC, Liliputing

↗

and Harpers LLC (2012 Kindle Fire). GeekWire

↗

Amazon also files under its own name, "Amazon.com Services LLC." AFTVnews

↗

- **Conflict to flag:** a third-party site (fccidlookup.com) describes "**Prime King LLC**" (Ridgeland, Mississippi) as "a known proxy for Amazon hardware" Fccidlookup

↗

for a Wi-Fi 6E/Zigbee Fire TV (FCC ID 2BFXH-1523). However, independent verification against FCC grantee databases could **not confirm** "Prime King LLC" or "Rocket Networks" as Amazon proxies. Treat Vlado L.L.C. as confirmed; treat Prime King LLC as **reported-but-unverified**.

- **Regulatory criticism:** The shell-filing practice is **legal** (the FCC permits an applicant/agent other than the true manufacturer) FCC ID

↗

and there is **no FCC enforcement action** against it. Criticism is journalistic (secrecy to thwart leaks), not regulatory. Separately, the FCC *has* criticized Amazon on adjacent grounds — e.g., a 2024 Commissioner letter over insecure video doorbells sold on Amazon lacking visible FCC IDs, and an FCC investigation into Amazon's marketing of signal jammers.

6. FTC AI-accuracy guidelines applied to hardware-ID errors

- **Foundational guidance:** FTC Business Blog, "Keep your AI claims in check" (Feb 27, 2023, Michael Atleson): performance claims must be substantiated, not "science fiction"; you "can't just blame a third-party developer"; the "black box" defense is rejected. Federal Trade Commission
- ↗
- **Enforcement trajectory:** "Operation AI Comply" (launched Sept 2024) → 2025 brought a documented wave of "AI-washing" actions under **Section 5 of the FTC Act** (no AI-specific statute needed):
 - **DoNotPay** — final order approved by a unanimous 5-0 Commission vote on January 16, 2025: the company must "pay \$193,000 in monetary relief" for deceptive "world's first robot lawyer" claims and notify 2021-2023 subscribers.
 - **Workado / Content at Scale** — complaint and Decision and Order finalized April 24, 2025 (matter 232-3092): Workado's "98.3% accuracy" AI-detector claim was unsubstantiated because the model — built by Norwegian students and trained only on scholarly-article abstracts — correctly detected AI-generated non-academic text "merely 53.2% of the time," with Director Chris Mufarrige noting "the product did no better than a coin toss."
 - **Evolv Technologies (2024)** — AI threat-detection overclaims.

- **Application to your context:** The throughline is *substantiation of accuracy claims*. An AI assistant or security tool that asserts a specific chip identity (e.g., “this device uses DRAM X” when it is actually an eMCP) is making an unsubstantiated efficacy/accuracy claim. The FTC framework does not yet have a hardware-identification case on point, but the doctrine (Section 5; competent-and-reliable-evidence standard; no black-box defense) extends naturally. **State this as an extension of existing doctrine, not as a settled FTC position on hardware misID.**

7. Prior art on “poisoned training corpus” via FCC documentation gaps

- **The exact framing — FCC §0.459 redactions creating training-corpus gaps that cause AI hardware-misID — does not appear in the literature.** This is the novel contribution of your report.
- **But every component is established:**
 - **Data poisoning:** Alber et al., *Nature Medicine*, Vol. 31 (Feb 2025; DOI 10.1038/s41591-024-03445-1) found that “replacing just one million of 100 billion training tokens (0.001%) with vaccine misinformation led to a 4.8% increase in harmful content ($P = 0.03836$), achieved by injecting 2,000 malicious articles ... that we generated for just US\$5.00,” and that a biomedical knowledge-graph defense “captures 91.9% of harmful content ($F1 = 85.7\%$).” Surveys (arXiv 2405.12750, 2506.23762) treat training-data poisoning and “incidental data poisoning” via outdated/misleading web content as known threats.
 - **Long-tail / knowledge-gap hallucination:** multiple surveys (arXiv 2311.05232; 2510.06265) establish that LLMs fabricate on rare entities and obscure part numbers underrepresented in training data — precisely the failure mode for a chip whose authoritative record was redacted.
 - **Documentation-gap → fabrication:** well-documented for knowledge-cutoff and niche domains.
 - **Your defensible thesis:** redaction under §0.459 plus the “already-public-isn’t-confidential” asymmetry means the *correct* chip identity is systematically absent from the scraped FCC corpus, while *incorrect* distributor/aggregator descriptions (which mislabel eMCP as DRAM) are present — yielding a structurally biased, “poisoned-by-omission” training signal that propagates hardware misidentification. Label this clearly as an original hypothesis bridging known primitives.

Details

Why the technical core is solid. The RPMB attack literature is recent, peer-reviewed/preprinted, and reproducible: Fukami & Buurke’s EMFI work bypasses RPMB authentication on multiple eMMCs (Samsung controllers) with a sub-\$7,000 rig, and the companion key-extraction paper shows the RPMB key can be pulled directly from the chip and used to defeat anti-rollback and undo data wipes. Combined with CERT VU#231329 (RPMB replay weaknesses), this establishes that the security guarantees of an eMCP like the H9TQ64A8GTDC depend heavily on correct identification and correct TEE provisioning — which is exactly why a misidentification (eMCP-as-DRAM) is a security-relevant error, not a pedantic one. If a defender or an AI tool believes a device has no managed NAND/RPMB when it does, the entire anti-rollback/secure-storage threat model is mis-scoped.

Why the FCC/AI bridge is the report’s original contribution. The FCC’s own rules create the asymmetry: BOM/schematics/internal photos are routinely withheld (permanent or 180-day short-term confidentiality), while the “information already public is not confidential” doctrine means the corpus skews toward whatever non-authoritative descriptions already circulate (distributor pages that flatten eMCP to “DRAM,” forum posts, etc.). The data-poisoning and long-tail-hallucination literatures show that such skew is sufficient to produce confident, wrong outputs. The FTC’s Section 5 substantiation doctrine supplies the accountability hook. No source ties these together for hardware-ID specifically — that synthesis is yours.

Caution on speculative/forward-looking material encountered. Several sources were promotional or predictive and should not be cited as fact: the fccidlookup.com Fire TV writeup speculates about launch timing (“strongly points to a product launch in Q4 2025”) and asserts Prime King LLC is an Amazon proxy without primary FCC corroboration; vendor pages (Kingston, BIWIN, Lexar) are marketing; LLM-hallucination “statistics” aggregator pages (e.g., sqmagazine) cite unverified percentages. Use the peer-reviewed/primary sources instead.

Recommendations

Stage 1 — Package the report for Krebs’s intake (do first).

- Send a concise pitch to krebsonsecurity@gmail.com AND request to continue on Signal `briankrebs.07` for anything sensitive. Lead with the consumer-impact hook and the single most newsworthy, verifiable fact.
- Lead the pitch with named, verifiable specifics — FCC ID 2AE6S-0948 / Vlado L.L.C. / Fire TV Stick LY73PR, and the H9TQ64A8GTDCUR-KUM part number — to clear the spam filter Krebs describes.
- Include primary exhibits: FCC filing screenshots, the SK Hynix part-number decode, chip photos, and citations to Fukami/Buurke and CERT VU#231329.

Stage 2 — Separate fact from hypothesis explicitly.

- Present three confirmed pillars (eMCP/RPMB attack surface; FCC §0.459 redaction mechanics; Amazon shell-entity filings) as established.
- Present the “poisoned-corpus-by-omission → AI hardware misID” as a clearly labeled original hypothesis, supported by the data-poisoning and hallucination literature, not as proven.
- For any AI-accuracy/FTC angle, frame as a doctrinal *extension* (Section 5 substantiation) rather than existing enforcement on hardware misID.
- Be scrupulous that the EMFI study used Samsung eMMC controllers, not SK Hynix — cite it as evidence of the *device-class* attack surface, not as a result on the H9TQ64A8GTDC.

Stage 3 — Pre-empt the obvious rebuttals.

- Get Amazon and SK Hynix on the record (Krebs will expect this). Note Amazon’s FCC compliance contact is Amazon.com Services LLC, 410 Terry Ave N, Seattle. [Amazon](#)
↗
- Address the “shell filings are legal” point head-on; your critique is about transparency/training-data consequences, not illegality.
- Reproduce or cite a reproducible demonstration of the misidentification (e.g., query a popular model with the part number and show the eMCP-as-DRAM error) to move the AI claim from assertion to evidence.

Benchmarks that would change the recommendation.

- If a CVE or vendor advisory is issued specifically for H9TQ64A8GTDC, elevate the technical severity framing.
- If “Prime King LLC” / “Rocket Networks” can be tied to an Amazon grantee record, add them; until then cite only Vlado L.L.C. and the corroborated shell names.
- If the FCC issues any AI-training or transparency guidance, or the FTC brings a hardware-ID substantiation case, upgrade those sections from “extension” to “on point.”

Caveats

- **Coverage gap is real:** no prior Krebs article on this exact topic was located; pitch accordingly but verify before claiming “Krebs has never covered X.”

- **Prime King LLC / Rocket Networks unverified** as Amazon proxies; only Vlado L.L.C. (2AE6S) is confirmed.
- **No PGP key** for Krebs was found; do not assume encrypted email is available — use Signal.
- **No part-specific CVE** for the H9TQ64A8GTDC; the device-class RPMB/eMMC research and the DDR5 Rowhammer CVE (CVE-2025-6202) are related but distinct, and the EMFI research targeted Samsung, not SK Hynix.
- **The FCC rules are silent on AI**; the §0.459-to-training-corpus causal chain is your argument, not regulatory fact.
- **FTC has no hardware-misID AI case**; the application is doctrinal extrapolation from Section 5 / Operation AI Comply.
- Some retrieved sources were marketing or speculative (launch-date predictions, aggregator hallucination stats); these were excluded from factual claims and should not be cited.