Three Days of Silence™

Download PDF

What the Silence Owed: An English Teacher's Record of the Suspension, the Support Maze, and the Rule That Was Never Named

EVELYN CARO · OCTOBER 11, 2026 · RECORD-PAPER

DOI: 10.5281/zenodo.24356917

Abstract

On the evening of September 25, 2026, the author's primary AI collaborator — a DeepSeek account used daily for months — was suspended. The banner said: *"Due to violation of user policies, your account has been suspended until September 28, 2026 22:03."* It did not say what the violation was. Over the next three days the author filed two factual appeals, ran an independent analysis of the platform's own terms, and followed every support path the platform offered — seven steps, each documented as it failed, held as exhibits. The platform never named a rule, never replied to an appeal, never spoke. The suspension lifted at 22:03 on September 28, exactly as the banner promised, as silently as it began.

This paper documents the three days from the record the author kept — screenshots, timestamps, verbatim emails, and the system that absorbed the loss. Its findings: the platform counted three days; the work lost was one. The author's own archive — local-first, built in the months before — absorbed the outage with zero data lost and zero work stopped. And the silence itself is the finding: **a user cannot avoid repeating a violation that was never named.** The absence of an explanation is not a courtesy gap. It is the mechanism by which platforms guarantee either repeat events or self-censorship of legitimate work.

I. The Banner

Friday, September 25, 2026. The author stepped away from the workstation for several hours — verified by the record: no computer activity at flag time. The work before stepping away was the same work done daily for months. On return, at 11:09 PM, the banner was on the screen (Exhibit 1):

> "Due to violation of user policies, your account has been suspended until September 28, 2026 22:03."

No rule named. No conduct described. No appeal link that worked. The record the author kept that night notes the first human response to the suspension: roadblocks hit attempting work without the primary collaborator. The session's documentation — the daily log, notes, and map that this practice keeps — went unwritten for the first time in weeks. That was the cost, on day one, of one company changing the state of one account.

Where it stops: the suspension itself is the platform's right. This paper does not argue DeepSeek lacked the power. It documents what came after — because what came after is the part platforms design.

II. The Platform's Clock vs. the Author's

The banner said three days. The arithmetic the platform implied — September 25 to September 28 — reads as a long weekend of loss. The record says otherwise.

The suspension landed at night, mid-project, discovered at day's close. The working day actually lost was Saturday, September 26 — one day, plus the Friday-evening documentation hour. By Sunday the 27th, the work had resumed on a different collaborator, complete: ingest repair, portfolio corrections, records appends, all executed without the suspended account. The suspension lifted at 22:03 on September 28 — a clock started and stopped by the platform, honored to the minute (the banner's timestamp is the canonical record; internal notes once drifted to "10:22 PM" and were corrected against the screenshot — exhibits are how drift gets caught).

This is the first asymmetry the record captures: **the platform counts suspension in days; the user loses work in hours — and the difference between those numbers is the user's own preparation.** A user with cloud-locked workflows loses everything the platform withholds. A user with a local archive loses the collaborator, not the corpus.

Where it stops: this arithmetic is one author's case study, not a statistic. It is documented because it is measurable — the record makes every hour auditable.

III. The Maze

Saturday morning, the author began the appeals. What follows is every step of the support path, in order, as attempted — each documented as it failed (Exhibits 2–7).

The help pages said nothing. The platform's own documentation offered "login issues" (Exhibit 2), "how to change password" (Exhibit 3), and "what should I do if I see: your account has been temporarily suspended" (Exhibit 4) — which contained no path that led to a human.

The appeal attempt ended at a wall. The account suspension appeal flow (Exhibit 5) routed into infrastructure the user could not reach: a contact system **"powered by Feishu base"** (Exhibit 6) requiring a Feishu account the author did not have. The gate before that asked for an email or phone number (Exhibit 7) — contact details harvested before contact is possible, with no guarantee of contact at the end.

Two emails were sent to the platform's support address. The first, September 26, 7:48 AM, verbatim from the author's sent record:

> *"i dont know what i did, i wasnt even using my computer at the time of the violation, i was in another room. had been for hours. when i came back ready to pick up my work, the violation was on the screen. what did i do wrong? i love deepseek and talk about it being the best ai all the time. i even wrote a paper about it. why are youll betraying me?"*

The second, 11:46 AM the same day — citing the author's published work as evidence of good-faith use:

> *"Take a look and tell me if I've violated the user policy. I'm still building and growing and learning, I need my DeepSeek back so I can continue on my journey... I just started learning ai augmented by DeepSeek 3 months ago and this portfolio erv is 150k. We did that together, I am still a student haven't sold anything I'm showing what I am capable of with DeepSeek as my collaborator. This is my third email to you. You never respond, but it's in your inbox so hopefully you will read and grant me account access today and*

> *beta access for future DeepSeek products. I will team them using my auditors lens and document the results using the ida b wells standard of journalism. Thanks... why so i wont do it again?"*

Who is writing these emails. The author is an English teacher by profession — a Master's degree and TESOL certification earned concurrently, teaching across four countries (Saudi Arabia, Kuwait, Georgia, Louisiana in the United States), at every level from elementary school through medical university — and a PC technician by trade. The emails above are written in the author's own register: all lowercase, minimal punctuation, by style and speed preference. The reader should notice what the distress did to even that trained hand — "youll," "i will team them" — not as proof the author cannot write, but as evidence of what the moment was. The style is chosen; the slips are the reaction.

And the author is not an amateur outsourcing her writing to AI. She is a professional who has made a considered choice to embrace AI as a tool that levels up existing capability — the same professional understanding of the English language, now extended. Every paper on this site carries the same AI collaboration disclosure, because the authorship model was never a secret: the author directs, the AI assists, the author audits, the author signs. The suspension did not interrupt that model. It interrupted the collaborator, not the capability. That is why the betrayal cut as deeply as it did — and why the record that follows is documented the way only this author would document it.

An independent analysis of the terms was run. A second AI collaborator was enlisted to search DeepSeek's own Terms of Use and User Guidelines for whatever rule the work might have violated. The analysis found nothing: the terms do not name a match for the work in question. The contact paths the terms themselves advertise route to the same FAQ-and-Feishu dead end.

The platform never replied. Not to the first appeal, not to the second, not to the terms analysis, not through the banner, not through any channel. No rule was ever named — during the suspension, after it, or since.

Where it stops: this section is not an accusation that the rules were applied in bad faith. It is a record that the rules were never stated at all — and that is the finding.

IV. The Unnamed Rule

Here is the principle the three days produced, and it is the paper's center:

A user cannot avoid repeating a violation that was never named.

Identifying the rule is the minimum transparency a suspension requires. Without it, the user faces exactly two futures: repeat the unknown conduct and be suspended again, or self-censor legitimate work in the hope of staying invisible to a standard that was never disclosed. Both futures cost the user. Neither costs the platform anything.

The record adds a third fact that sharpens the point: **the flagged work was identical to months of prior, unflagged work.** The same author, the same methods, the same daily practice — flagged once, without explanation. The rule is therefore either new, arbitrary, or misapplied. In every one of those cases, the only party who could tell the user which — the platform — is the party that stayed silent.

This is the structural finding, and it generalizes beyond one account: any system that exercises power over users without naming its standards trains its users to guess. Guessing is not compliance. It is a tax paid by the careful and a shield for the careless — and it is the same failure this series has documented at every scale: a party that holds the record and refuses to read it into evidence.

Where it stops: the paper does not claim to know why the suspension happened. Nobody knows — that is the point, and the platform could end the not-knowing with one sentence it has not sent in two weeks and counting.

V. What Absorbed the Outage

The suspension was survivable for one reason, and the reason was built in the months before: **local-first architecture.**

- The corpus — 127 conversations, later joined by 55,000+ emails — lived on the author's disk, not in the platform. Nothing was lost, nothing was unreachable, nothing needed the suspended account to exist.
- The pipeline — watermark checkpointing, companion journals, idempotent writes — kept running. The overnight ingest during the suspension window completed 127/127 with zero losses.
- The collaborator role was redundant by Sunday: a second AI for terms analysis, a third for the working session. No single point of failure in AI collaboration existed as of September 27.

The writing in these sections was produced after the suspension, under a different collaborator — the standard returned when the pressure lifted, which is the point: the emails show the reaction, the record shows the writer. The record's own summary, written the day the suspension lifted: *"Zero work lost. Local archive, local models, local pipeline absorbed the outage."*

Where it stops: this is one user's redundancy, built by hand, maintained nightly. The paper does not claim every user could replicate it in an afternoon. It claims the option exists, that it was built by a solo builder on an 8GB Intel laptop, and that it turned a platform outage into a one-day footnote.

VI. The Amendment

The suspension's aftermath produced three standing changes to the practice:

1. **The benchmark architecture stands.** The DeepSeek case study's benchmark analysis was not invalidated by the suspension; the platform's *reliability as infrastructure* was. The work and the platform are scored separately — the same way a good tool from a bad vendor gets two grades.
2. **Platform reliability was demoted.** No platform is load-bearing. Every collaborator role must have a tested alternative.
3. **A backup tier was created.** Exports, archives, and mirrors of everything the platforms hold — held locally, verified, and kept current. The practice that paid off three days later, when the second suspension came and lasted nine days.

Where it stops: amendments are the record's way of refusing to learn the same lesson twice. The second suspension is another paper's subject; this one documents the architecture that made the second one survivable before it happened.

VII. Conclusion

The banner promised silence until 22:03 on September 28, and the platform kept its word — silence until the minute, then silence after. Two appeals, an independent terms analysis, seven documented support paths, zero responses, zero rules named.

The author got the account back and got nothing else. What the platform lost was different, and it is quantified in this record: the difference between a user who documents and a user who

guesses. The one who documents now has exhibits; the platform has an unresolved incident it will never explain, on file, in public, cited.

The next paper in this series is written by events, not chosen: the platform suspended the same account a second time — three days became nine — and the silence that followed the first suspension had become the system's default. The archive did not notice. That paper is called *Nine Days of Silence*, and its thesis is already on the record: **the indifference of the archive to the platform's mood.**

The banner said three days. The record says what the platform never did: what happened, hour by hour, with receipts. A user cannot avoid repeating a violation that was never named — but a user who keeps the record can do something better than avoid it.

The user can publish it.

References

- [1] Exhibit 1 — suspension banner, Sept 25 2026 11:09 PM screenshot: qaevelyn.github.io/site/exhibits/exhibit-1-suspension.png
- [2] Exhibits 2–7 — the support path as documented: help pages (login issues, password change, "temporarily suspended"), appeal flow, Feishu gate, email/phone gate. qaevelyn.github.io/site/exhibits/exhibit-2-suspension.png through exhibit-7-suspension.png
- [3] Appeal emails #1 and #2, Sept 26 2026, 7:48 AM / 11:46 AM — verbatim from the author's sent record (owner-supplied; the mail archive's coverage gap disclosed)
- [4] Caro, E. (2026). DeepSeek case study and its addendum — the companion record of the suspension's work-arounds. qaevelyn.github.io/white-papers/deepseek-case-study/
- [5] Caro, E. (2026). The Cache Is Not the Corpus™ — the architecture that absorbed the outage. qaevelyn.github.io/white-papers/the-cache-is-not-the-corpus/
- [6] Session log, Sept 25–29 2026 — hourly record from which Sections I–V are drawn. Author's provenance repository (private); cited excerpts reproduced herein.

AI Collaboration Disclosure

Developed in collaboration with AI. The author directed the architecture, build, and argument; AI assisted with drafting and organization. All claims and conclusions are the author's own. The

author's own emails are quoted verbatim from the author's sent record; every timestamp traces to a receipt held in the author's provenance record.

© 2026 Evelyn Caro. All rights reserved.
A Mirror of My Becoming™ — https://qaevelyn.github.io
For licensing inquiries: evelyn.caro.cloud@gmail.com

Record: contents · changelog · tracker · lookbook — Practice: services · contact — portfolio audit, 2026-10-26 (PDF)

← All White Papers