

Re: Update on your request #00697534 [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

From: Soul Legion <legionkillfeed@outlook.com>
Date: Wed 2025-12-10 10:15 AM
To: Trust Support <trust.operations@trustpilot.com>
Cc: Trustpilot Legal <legal@trustpilot.com>

I am in **New South Wales, Australia**. The standard that matters here is **Australian law**, not whatever an internal policy or an automated tool is willing to let slide.

In particular, you need to consider:

- **Defamation Act 2005 (NSW)** – including the new **section 31A** on digital intermediaries and the traditional **section 32** defence of innocent dissemination. Once you receive a written complaint that clearly identifies the plaintiff, the matter and why it is defamatory, you only keep your digital-intermediary defence if you take **"reasonable access prevention steps"** within 7 days. [NSW Legislation](#)
- If you choose to leave the material up after being given clear, technical proof that it is false and cannot have occurred, you are no longer an innocent conduit. You are a primary publisher with knowledge. Every day you continue to host it and every time you refuse removal after a complaint is treated as a fresh publication.

I have already given you written complaints that:

1. Identify me and my business;
2. Identify the specific review and where it appears;
3. Explain in detail why the matter is defamatory and technically impossible.

That is exactly the type of complaint contemplated by section 31A(3) of the Defamation Act. [NSW Legislation](#)

Once that happens, you **do not** get to hide behind "our guidelines were applied" while doing nothing meaningful to prevent access. The law is not "we ran it past an AI and it said okay".

Regards,

Anthony Brodie
LEGION KILLFEED



From: Soul Legion <legionkillfeed@outlook.com>
Sent: Wednesday, 10 December 2025 9:30 AM
To: Trust Support <trust.operations@trustpilot.com>
Cc: Trustpilot Legal <legal@trustpilot.com>
Subject: Re: Update on your request #00697534 [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Subject: Re: Case #00713084 – month-on-month damage since August and your continuing refusal to act
Hi Nicole,

It has now been almost five hours since my last email. For someone presented as a "Senior Escalation Specialist" dealing with a review your own systems have auto-flagged as potentially harmful or illegal, that silence speaks volumes. A real specialist would not sit on an auto-flagged, legally sensitive review like this while a small business owner is begging you to stop publishing something you know is false.

I want the following on the record.

1. Month-on-month damage since August

You have asserted that it is "improbable" that content like this would cause the level of serious harm required for defamation. My internal data says otherwise.

Looking only at active paying subscribers (unique paying accounts) and ignoring all other revenue sources:

- From early 2024 up to August 2025, my subscriber numbers steadily climbed into the high double-digits, just under the 120 mark.
- From August through November 2025, there is no month in which the trend reverses or stabilises. Each billing cycle after August records fewer active subscribers than the one before.
- By the end of November 2025, I am down to roughly three-quarters of the August subscriber base. In other words, around a quarter of my paying subscribers have vanished in four months, with no organic recovery.

That is not a one-off wobble or normal churn. It is a clear month-on-month decline that starts once the organised review and harassment campaigns fired up again and accelerates through the period where you decided to host and defend this "Hooded_saw / Consumer" review and the other negative reviews tied to the same dispute.

You do not get to wave that away as "improbable" just because acknowledging it would be inconvenient. You also do not get access to my raw financials simply because Content Integrity wants to pre-argue its defence by email. The underlying records will be placed before my lawyers, a judge and regulators, not your triage queue. Your dismissive assertion that serious harm is unlikely will then be weighed against hard numbers and the obvious timing.

2. The November 30 cluster – including a competitor

You also keep pretending this is just one angry customer. It is not.

On 30 November 2025 you received, in very close succession:

- the edited "Consumer / Hooded_saw" review;
- a review from "Kamy", a documented direct competitor; and
- a review from "Unholy", part of the same network of hostile actors.

All three land on the same date, all draw on the same manufactured narrative that I mistreat customers, and all tie back to the same incident:

- A non-customer ("Hooded_saw") trying to bypass my support process to access someone else's server;
- the real subscriber, lih0u2451, never once stating that this person was an authorised representative or that they had any right to speak on their behalf; and
- the moment I refused to treat that stranger as the account holder and called out the behaviour, your platform suddenly fills up with negative content – including from a competitor whose own commercial interest is obvious.

Yet you still insist you "cannot arbitrate" the facts and are forced to assume these are genuine, independent experiences.

What you are actually doing is giving hostile third parties and competitors a public megaphone, then hiding behind your internal processes when the target of that attack points out the obvious coordination.

3. Failure to respond and the question of whether this was a real specialist

You signed off as a "Senior Escalation Specialist" and then:

- asserted that the damage was "improbable";
- implied I was effectively wasting your time by treating this as defamation; and
- disappeared for hours without addressing the very specific legal and structural issues I raised.

That behaviour, combined with the boilerplate tone of your email, makes me genuinely question whether I am corresponding with an actual specialist at all or just another template being pushed through a queue.

Either way, your failure to act on a review that your own system has auto-flagged as potentially harmful or illegal, in a case you know involves a concentrated campaign by competitors, is not what good-faith escalation looks like. It is what liability management looks like.

4. Your systems do not override Australian law

Once again: your internal "high bar", your sense of what is "unlikely" to be harmful, and your content-integrity checklists are legally irrelevant.

Australian defamation law does not say "Trustpilot may keep publishing whatever it likes provided a Senior Escalation Specialist has ticked all the boxes."

What matters is:

- that you continue to publish and republish this review into Australia;
- that you have actual notice that key factual claims in it are false (no subscription, no ticketless customer, no authorised link between the actual subscriber and the reviewer);
- that the review forms part of a cluster of hostile reviews, including competitor reviews, on the same day; and
- that you are now on notice of a clear month-on-month collapse in my subscriber base following that cluster.

Your insistence that, according to Trustpilot's internal worldview, this is all unlikely to be seriously harmful does not change those facts. It simply increases your exposure when a court comes to a different conclusion.

5. Brand manipulation and the fiction of "Trust"

You are doing all of this while trading under the name "Trustpilot" and marketing yourselves as a neutral, reliable arbiter of "genuine customer experiences".

Continuing to host:

- an edited, auto-flagged review that your own system marks as potentially harmful or illegal;
- competitor reviews tied to the same manufactured dispute; and
- a misleading "transparency" page that paints me as the problem while hiding your repeated missteps
- is the opposite of trustworthy behaviour.

Your conduct misleads the public into believing that content on your platform has been meaningfully vetted and can be relied upon. In reality, you are knowingly keeping false, coordinated and auto-flagged material online while insisting that your brand and internal policies somehow put you above the law that applies to every other publisher.

6. The personal impact you keep ignoring

Finally, I need to put this on the record because you clearly have no intention of acknowledging it otherwise.

This situation, across Trustpilot, Discord, Reddit and other platforms, has dragged on for years. I have explicitly told Trustpilot that the ongoing defamation and harassment – much of it amplified by your refusal to act – has pushed me to the point of suicidal ideation. That was not hyperbole. It was an honest statement of how severe the impact has been on me personally.

Nobody at Trustpilot has ever:

- acknowledged that statement;
- expressed even basic concern; or
- indicated that it would be taken into account in how you handle my case.

Instead, your so-called specialist wrote to tell me that my concerns are unlikely to meet your internal threshold for serious harm, while you continue to publish content your own system has labelled as potentially harmful or illegal.

To be clear: that pattern – ignoring explicit statements of suicidal distress while refusing to remove debatable content and then lecturing me about your "high bar" – will be relied on as evidence of bad faith in any proceedings. It goes directly to aggravated damages and to the question of whether Trustpilot has behaved like a responsible, neutral intermediary or as a company more interested in defending its PR line than protecting people it knows are being harmed.

7. Next steps

- Your position is now very simple:
- You are fully aware that the "Consumer / Hooded_saw" review was written by a non-customer using an Apple relay address, then heavily edited.
- You know it sits in a cluster with competitor and associate reviews posted on the same date.
- You know that, from August onwards, my subscriber base has fallen month after month and that by November I had lost roughly a quarter of my paying customers.
- You know your own system has auto-flagged that review as potentially harmful or illegal.
- You are still choosing to keep it online and to lecture me about your internal thresholds.

If that is truly the stance of a "Senior Escalation Specialist", then I will treat it as Trustpilot's formal position and proceed on the basis that:

- this correspondence;
- your failure to act within hours of receiving detailed notice; and
- your dismissal of obvious harm

will all be put before a court and the relevant regulators as part of a pattern of knowing, reckless publication.

I again invite you – or someone with actual legal authority inside Trustpilot – to reconsider, remove the review, correct the record, and start behaving in a way that matches the word "Trust" in your company's name.

Regards,

Anthony Brodie

**LEGION
KILLFEED**



From: Soul Legion <legionkillfeed@outlook.com>

Sent: Wednesday, 10 December 2025 4:50 AM

To: Trust Support <trust.operations@trustpilot.com>

Cc: Trustpilot Legal <legal@trustpilot.com>

Subject: Re: Update on your request #00697534 [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Hi Nicole,

How dare you imply that the harm I am complaining about is trivial, "improbable", or essentially frivolous.

You are looking at a situation where:

- your own system has placed a banner over the "Consumer / Hooded_saw" review saying it may contain harmful or illegal content;
- you know that the review was heavily edited after your original investigation;
- you know that the edited version now contains detailed factual allegations that are flatly inconsistent with my records; and
- you have been told for over a year about a broader pattern of review-bombing and competitor harassment tied to exactly these reviews.

To then tell me that this sort of publication is unlikely to cause serious harm is not a neutral assessment – it is a position you are taking to protect yourselves.

1. Your internal "high bar" is irrelevant to Australian defamation law

You keep repeating that Trustpilot has an "exceptionally rigorous threshold" for defamation and that you only remove "the most serious and consequential" content. That might be your **internal policy**, but it is completely irrelevant to the **Defamation Act 2005 (NSW)** and the serious harm test that actually applies to publications in Australia.

The law does **not** say:

- "platforms may leave content up so long as they personally think it is unlikely to cause serious harm"; or
- "global review sites can invent their own stricter standard and ignore the statutory one".

Once you choose to publish and republish reviews accessible in Australia, your obligations are set by **Australian courts and legislation**, not by Trustpilot's marketing language. Whether "most" reviews on your platform are harmless is beside the point. The only question that matters is whether **this** review, in **this** context, has caused or is likely to cause serious harm to my reputation and business.

That is a fact question that will ultimately be determined by a judge, on proper evidence – not by Content Integrity deciding that my losses are "improbable" because it would be inconvenient for you if they were otherwise.

2. This is not "limited commentary" – it is a detailed story you know is false

You describe the content as if it were a bit of rude language or a single negative word. That is not what you are publishing at:

- <https://www.trustpilot.com/reviews/692b14c5772dcebff6caec38>
- <https://www.trustpilot.com/review/killfeed.co>

The edited review now online contains a long narrative, including assertions that:

- the reviewer was a **paying customer**;
- I **refused to help** because they "weren't worth my time";
- they "made a ticket but got no response";
- I created a **separate channel** to call them a moron, imbecile and idiot; and
- I was **sharing screenshots** to "show off" to my community.

Those are not vague impressions. They are concrete statements about whether someone bought my service, whether I provided support, and how I behave towards customers in internal channels.

You know – because I have already sent it to you in other cases – that my subscriber and ticketing records tell a completely different story:

- the only email you are treating as "verified" is an Apple relay address;
- there is no subscription under that address or under the name "Hooded_saw";
- the actual paying subscriber was another person entirely, who **did** receive multiple responses but chose to ignore them; and
- "Hooded" was a third party trying to bypass the ticket process, threaten to move funds, and weaponise your platform when I refused to play along.

You cannot honestly call that "limited commentary" or "a solitary negative word". It is a detailed factual story that you have every reason to know is untrue.

3. Clustered negative reviews and the timing of subscriber drop-offs

You also claim that it is "improbable" that this type of review would cause serious financial or reputational damage. That statement is completely disconnected from the reality of a small subscription business whose entire customer base overlaps with your user pool.

Without going into specific figures (which I am deliberately reserving for court and for regulators):

- For a long period, my subscriber numbers **grew steadily** despite the occasional one-star review. Single, isolated negatives in 2023 and early 2024 did not coincide with any downturn.
- In mid-2025 there is a **cluster** of one-star Trustpilot reviews over a short window, including content from direct competitors and from this same "Hooded_saw" incident.
- Immediately after that cluster, new sign-ups slow dramatically and cancellations increase. Instead of continuing to grow, my paid subscriber base **turns and starts falling**, and that decline persists across multiple billing cycles.

In other words: isolated negative commentary did not visibly affect growth, but a concentrated burst of hostile reviews – including the one you are now defending – correlates with a clear change in direction. You may choose to argue in court that this is coincidence, but you do not get to brush it away as "improbable" with no evidence.

You **do not** have my full subscription and financial data, and you **will not** be getting it just because Content Integrity would like to eyeball it. That material will be produced where it belongs: to my lawyers, to a judge, and if necessary to regulators. When that happens, your statement in this email that serious harm is "unlikely" will be

compared against hard numbers and an obvious timeline.

4. You cannot call yourself neutral while branding this as “verified”

You say you “cannot definitively determine the objective truth” of what happened and must remain neutral in factual disputes. At the same time, you:

- label this as a **review of my business**;
- display it as a **verified** experience;
- in the past, have even shown it as if it were **invited**; and
- use your “transparency” page to paint me as a problem business that keeps flagging honest reviews.

Those are not neutral acts. They are positive representations to the public that:

- this person genuinely paid for and used my service; and
- their description has been vetted through your integrity systems.

You cannot hide behind “we don’t know who is right” while still stamping one side of the story with a Trustpilot badge and pushing it to anyone who searches my name.

If you truly believed you could not adjudicate the facts, the bare minimum would be to **remove the verified/invited status**, or take the review down until the underlying relationship is clarified. Instead, you have chosen to keep the most damaging framing while claiming your hands are tied.

5. Auto-flagged as harmful/illegal, edited after investigation, and I am blocked from re-flagging

The most damning part of this is your **own banner** on the direct review URL, stating that your automated system has flagged the review as potentially “harmful or illegal”.

At the same time:

- the review remains fully visible;
- it is still presented as verified; and
- the business being targeted is told “You’ve already flagged this review”, and is **locked out** of re-flagging, even though the version you investigated and the version now online are not the same.

You have been repeatedly told that:

- the original review was short and vague;
- the current version is a much longer edited text; and
- all your “investigation complete – no breach” decisions relate to the **pre-edit** version.

In practice, that means your workflow does this:

1. Investigate a short review.
2. Mark it as compliant.
3. Allow the reviewer to completely rewrite it via the edit feature.
4. Treat the new, longer version as already cleared.
5. Prevent the business from flagging the new content.

That is not an accident. That is a **design choice** that allows edited material to slip past checks while you continue to hold up “investigated and within guidelines” as a shield.

When that edited content is then auto-flagged by your own system as potentially harmful or illegal and you still leave it up, you are not a neutral intermediary. You are an active publisher with actual knowledge, continuing publication in full awareness of the risk.

6. Your assessment of harm will be relied on in court

You have now put in writing that:

- this kind of review “does not contravene” your harmful/illegal policy;
- you regard such content as **unlikely** to cause the level of serious harm required for defamation; and
- you do not consider removal “warranted at this time”.

You are entitled to that view. But you must understand what it means:

- If you keep this review online, after repeated notice and after your own system flagged it, each further day of publication occurs with **actual knowledge** of my objections and the broader harassment pattern.
- If proceedings are issued, this email will sit in the pleadings as evidence that Trustpilot was expressly warned about serious harm, chose to dismiss it as “improbable”, and continued to publish anyway.
- Your insistence that your internal threshold and neutrality language trump external law will go straight to aggravated damages and to any argument about reasonableness.

I am not going to hand you my subscriber ledger so you can pre-litigate your defence in Content Integrity. But I can and will put full financial, technical and contextual evidence in front of a judge. At that point, your statement that serious harm was “unlikely” will be one more piece of proof that you preferred your own model of the world to the reality in front of you.

7. What I expect from Trustpilot at this point

For this review and case, I now expect:

1. **Immediate and permanent removal** of the "Consumer / Hooded_saw" review in all its versions from public view.
2. A **clear explanation** of how edited reviews are handled, and why the current, heavily edited text is treated as if it had passed the same checks as the original.
3. Restoration of my ability to **flag materially edited reviews**, instead of being locked out because a different version was previously investigated.
4. Corrections to your **transparency page** so it accurately reflects the history of edits, flags, competitor involvement and late removals, rather than framing me as the problem.
5. Written confirmation that **all logs, edit histories, internal notes and communications** relating to this review, this case and the broader pattern of review-bombing will be preserved in full for litigation and regulatory review.

If you refuse to do that, then please at least be honest, on the record, that Trustpilot has:

- chosen to keep publishing a review your own system has flagged as potentially harmful or illegal;
- done so knowing that the "verified" status is not supported by payment records; and
- done so after being told that there is evidence of a clear temporal link between the clustering of hostile reviews on your platform and a subsequent collapse in growth for a small, subscription-based business.

That is the position you will be inviting an Australian court to examine.

Polite Regards,

From the increasingly harmed by Trustpilot's actions

Anthony Brodie - Developer of Legion Killfeed

From: Trust Support <trust.operations@trustpilot.com>

Sent: Wednesday, 10 December 2025 3:53 AM

To: legionkillfeed@outlook.com <legionkillfeed@outlook.com>

Subject: Update on your request #00697534 [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Hello Anthony,

Thank you for your reply. My name is Nicole and I am a Senior Escalation Specialist on the Content Integrity team.

While we understand you're frustrated, please note that we do not accept any threatening or abusive messages. I am here to explain our stance and why **this review will not be moved offline**.

Please keep in mind that we don't remove content just because you dislike or disagree with it or it criticizes your business. However, we want everyone to feel welcome on our site, and we don't allow any content that includes hate speech, discrimination, threats or violence, obscenities, defamation, or is terrorism-related.

After further investigation and assessment of the edited review, we confirm this does not contravene our guidelines for being harmful or illegal, and specifically for containing hate speech or discriminatory content. When assessing for hate speech and discrimination, we will remove content that negatively targets people or a group of people based on who they are, for example their religion, ethnicity, nationality, race, color, descent, gender, or other identity factor.

We acknowledge and understand your belief that this content is defamatory and harmful to your business. However, we must be clear: Trustpilot maintains an exceptionally rigorous threshold for removing statements on the grounds of defamation. This high bar is fundamental to our operation, as it ensures we only intervene in cases where the content is demonstrably and severely harmful. Our guidelines require the removal of only the most serious and consequential content.

Our process dictates that we analyze the full context of the review, rather than isolating single keywords or short phrases. In the vast majority of consumer-submitted online reviews—which are generally understood to be expressions of personal opinion—the use of a solitary negative word or brief critical comment is unlikely to cross the legal threshold for defamation.

It is improbable that such limited commentary, in this format, would be found to cause the level of significant financial loss or severe reputational damage necessary to warrant removal under our current policies.

We understand and acknowledge your profound disagreement with the reviewer's description of events. However, as a platform facilitating genuine consumer experiences, we are not positioned to arbitrate factual disagreements. We cannot definitively determine the objective truth of the claims made by either party. Consequently, we are compelled to take a strictly neutral stance in all factual disputes between businesses and reviewers.

Based on the current guidelines and our assessment, the review does not contravene our policies for being Harmful or Illegal and therefore, removal is not warranted at this time.

We understand that you disagree with the decision we've made, however, I can assure you it was taken after careful consideration and in accordance with our guidelines.

I hope this clarifies our stance, but if you have further questions, please let us know.

Best regards,

Nicole



See our reviews



[Visit Trustpilot](#) | [Help Center](#) | [Trustpilot Privacy Policy](#)

----- Original Message -----

From: Soul Legion [legionkillfeed@outlook.com]

Sent: 12/4/2025 1:00 PM

To: legal@trustpilot.com; trust.operations@trustpilot.com; press@trustpilot.com

Subject: Re: [thread:HUNjgfNQfSqlVgYTI2NNAGA::]

Subject: Case 00697534 – Notice of continuing publication of disability-based defamation (Hooded_saw review)

Hi Megha,

Thank you for your recent emails confirming that case **00697534** is considered "closed" and that the review:

"Rude, unwilling to help, Soul needs professional help."

Reference: **Hooded_saw**

will remain online.

This reply sets out:

1. Why that wording is disability-based harassment and plainly defamatory.
2. The full factual chronology from my Discord server (ticket, public chat and DMs), including timestamps and server status.
3. How your moderation history makes Trustpilot a knowing publisher.
4. What I am now requiring from Trustpilot.

For context I refer to three attachments I have already prepared:

- Attachment A: 2025-11-29-lihou2451-PUBLIC.png (public chat).
- Attachment B: 2025-11-29-lihou2451-TICKET.png (support ticket, screenshot captured 5 December 2025, 1:53am AEDT).
- Attachment C: 2025-11-29-lihou2451-hooded.png (DM with "Hooded_saw").

1. Why this is disability-based harassment, not a normal "opinion" review

In ordinary English, telling someone they "need professional help" is widely understood as claiming they are mentally ill or unstable and require psychiatric or psychological treatment. That is not a comment about product features or response times. It is a statement about my mental health.

Under Australian law, mental illness, and even an imputed mental illness, is treated as a form of disability. Using mental health as an insult in the context of providing services to Australian users is disability-based harassment, not harmless "opinion".

Your platform is not just reporting neutral customer feedback. You are publishing a sentence that combines:

- a mental health slur ("Soul needs professional help"), with
- assertions about my conduct as a provider ("Rude, unwilling to help"),

and attaching it directly to my business listing.

2. Why the review is plainly defamatory

Australian defamation law asks whether words would tend to:

- lower a person in the estimation of ordinary members of the community, or
- cause others to shun or avoid them.

Attached to my Trustpilot profile, the review:

"Rude, unwilling to help, Soul needs professional help."

clearly carries these imputations:

- I am professionally incompetent and deliberately unwilling to assist customers; and
- I am mentally unwell or unstable and therefore unfit to operate my business.

These are serious, factual sounding allegations about both my competence and my sanity. They are presented to potential customers as a "trusted" review of a real business.

As the chronology below shows, there is no factual basis for these allegations in the actual interactions. They are neither substantially true nor honest opinion based on proper material. They are simply abusive and misleading.

3. The actual facts: ticket, public chat and DMs

3.1 Who the real customer is (and is not)

- The paying subscriber is **Z Company**, whose Discord account is "**lihou2451**".
- "**Hooded_saw**" (**Frank**) is not the subscriber. He describes himself as an "admin" and friend. He has never been verified as associated with the account holder, never opened his own support ticket, and the account holder never responded to confirm him.
- My support terms are clear: I deal with the **ticket owner** for a subscription, not unverified third parties, so that there is an auditable support history.

Trustpilot is currently treating a non-customer's retaliatory comment as if it were a straightforward customer review.

3.2 My support system and the no-DM rule

In my Discord support server (see Attachment B):

- All assistance is handled through structured **ticket channels** created by my bot.
- The support panel explicitly instructs users **not to DM staff or my personal account** for support.
- This keeps all support in one logged place, avoids disputes about what was said, and ensures consistent treatment for all customers.

Despite that, the friend repeatedly ignored this and tried to conduct the dispute via DMs instead of using the ticket like everyone else.

3.3 The ticket: server running, customer not responding

(Attachment B - ticket; screenshot taken 5 December 2025, 1:53am AEDT)

On **29 November 2025**, my bot opened **#lihou2451-ticket** for Z Company. The bot and I:

- asked for a detailed description of the issue,
- requested screenshots or error messages, and
- asked which server it related to.

My diagnostic command in that ticket shows:

- "Z Company Server is online", and
- all feeds configured and responding.

Immediately under that, my message asks:

"Do you have logs where you can show this?"

The embed below that message shows:

- "Last Processed Line: 18:22:53 | ####", and
- "Last Msg Sent at 6 days ago" with a hover that displays the exact datetime "**Saturday 29 November 2025 5:33 AM**".

To be precise:

- The screenshot was captured on **5 December 2025**, which is why the Discord UI shows "6 days ago".
- The hover text in the screenshot confirms that my last ticket message was at **5:33am AEDT on 29 November 2025**.

By the time the later public chat messages occur that morning:

- I have already posted a detailed, logged response in their ticket at 5:33am, asking for logs and providing the last processed line.
- There is still **no reply at all** from the customer to those questions.

3.4 The public chat complaints instead of using the ticket

(Attachment A - public chat, same morning 29 November 2025)

Instead of replying in the ticket, **lihou2451** comes into my **public** channel and starts complaining that they have "lost the kill feed and hit feed again" and "just need a little help". The log in Attachment A shows that I:

- tell them directly that they **have not responded in the ticket**,
- paste a screenshot of that ticket showing my last unanswered message and the support embed, and
- repeatedly tell them to "just use the ticket" so I can help.

Crucially, in that public chat I **hyperlink the exact ticket channel** - the blue text "**#lihou2451-ticket**" - so all they need to do is click that blue link and type a reply in their own ticket.

They still do not do so.

Instead, they continue posting vague apologies and "I just wanted some assistance" messages in public, while still refusing to answer the specific questions that are sitting in their ticket.

This is happening at around **5 to 6am AEDT**. I had already been awake through the night working and supporting other customers. At that point my choices were:

- stay awake indefinitely moderating someone who refuses to use the proper channel, or
- temporarily stop the public disruption so I can sleep, while leaving the ticket open for a proper reply.

So, as shown in Attachment A:

- I **temporarily mute li hou2451 for one hour in public chat**, explicitly telling him "I'm gonna mute you for an hour and go back to sleep."
- The ticket itself remains open and unchanged, waiting for a substantive reply from him.

He is **not banned**, and I remain fully prepared to resume support as soon as he cooperates with the ticket process like every other customer.

3.5 The friend's DM, the ban, and when I replied

(Attachment C - DM)

Only after all of the above does the friend "**Hooded_saw**" come into the picture.

Attachment C shows:

- At **7:27am on 29 November 2025**, he sends a long DM to me introducing himself as "Frank" from the Netherlands, claiming to be an admin for Z Company, and saying his colleague is a "70-year-old British combat veteran" who loves DayZ.
- In that DM he **admits** that they "overlooked your request for clear information" and acknowledges that I cannot perform my role without the necessary details.
- He then says that if we "are unable to work together" they may move their server funds to a competitor.

At that stage:

- I do **not** reply to him.
- I focus instead on trying to get the actual subscriber, **lihou2451**, to respond in his ticket so I can even verify whether this "admin" is authorised to speak for the account.
- I still receive **no ticket reply at all** from li hou and no confirmation that the friend is staff.

Because the friend is:

- a non-customer,
- repeatedly ignoring the clearly posted no-DM rule, and
- trying to escalate privately while I still cannot verify his authority,

I ultimately ban "**Hooded_saw**" from my Discord **without engaging him in a DM conversation**. That ban is for rule breaches and disruption, not for asking for legitimate help.

Later, after all of this, and after **lihou2451 has left my Discord and cancelled his subscription**, I send **one blunt DM reply** to the friend - the message from me in Attachment C dated **30 November 2025 at 12:23am** - explaining that I cannot support people who repeatedly ignore the ticket system, choose public drama instead, and then send abuse.

Only after my customer has already left and the business relationship is over do I speak to the friend directly in DMs. His immediate response is:

"You are one sorry ass excuse
I'm laughing
Now block me"

That is the **end** of contact with him.

So the correct chronology is:

- I did not engage the friend before banning him.
- I tried repeatedly to get my actual customer to respond in his ticket so I could verify who this "admin" even was.
- Only once they had already left did I send a single DM explaining why I would not be continuing support.

3.6 Subscription status and server runtime

A few more concrete facts:

- At the time of the dispute the **Z Company server is online and functioning**.

- The support embed in the ticket also shows the server's **remaining runtime**. At the time I prepared this evidence bundle the server has a little under **two days of runtime** left in its Nitrado billing period. In other words, they have had the benefit of the service for the period they paid for.
- After the abuse from the friend, I tell **lihou2451** in the ticket that I will **still continue their subscription** if he can respond appropriately within 24 hours with an apology and a proper reply in the ticket.
- I do not cancel his subscription at that point. The choice is left with him.
- Instead of replying or apologising, **lihou2451 leaves my Discord server without any response**. Any cancellation after that comes from their side, not mine.

Throughout the time they remained in my server, I was willing to help, on the simple condition that they use the same ticket system as every other customer.

3.7 Pattern of only using public communication

When you put everything together, a clear pattern appears:

- When invited to respond privately in a logged ticket, **they refuse**.
- Even when I hyperlink their ticket in blue text so it is literally a one click action, **they still refuse**.
- Instead, they prioritise public channels that create visible "evidence" of their complaints:
 - my public Discord chat; and
 - your Trustpilot review platform.

Given the broader harassment campaign I have already documented across Discord, Facebook, Reddit and Trustpilot, it looks very much as if they were primarily interested in **public validation of their grievances**, not fixing any genuine technical problem. If the complaint had truly been about getting their feed working, the simplest possible step would have been to reply once in the ticket that was opened specifically for them.

The way the friend leans on the story of a "70-year-old British combat veteran" also raises a concern. It is unusual for people in that age bracket to be personally renting and configuring private DayZ servers. If li hou is genuinely unable to click his own blue ticket link or type a basic reply, it suggests that the younger "friend" may in fact be the one controlling the account and the complaints, while using the older person's age to add emotional weight.

Either way, the behaviour shown in the logs is not that of a helpless customer ignored by a "rude" provider. It is the behaviour of people who:

- ignore the support system that is there to help them,
- break clear rules about DMs,
- become abusive when boundaries are enforced, and
- then use Trustpilot to retaliate with a disability-based slur.

4. Your moderation history makes Trustpilot a knowing publisher

For this single review you have:

- rejected a flag that it constituted **hate speech or discrimination**,
- rejected a flag that it was **about a different business**,
- stated repeatedly that you will not consider legal arguments about defamation or discrimination and will only act based on your internal guideline categories, and
- after receiving the evidence above, told me again that the review is compliant and that no further response will be given.

At the same time you now hold:

- detailed explanations that the wording is alleged to be **defamatory and disability-based harassment** under Australian law,
- concrete evidence from my Discord server (ticket, public chat, DMs) showing that:
 - the reviewer is a **non-customer**,
 - the "unwilling to help" claim is contradicted by the ticket history, and
 - the "needs professional help" phrase has no factual basis and is purely an insult about mental health, and
- my formal notices that this review sits inside one of the most comprehensively documented multi platform harassment dossiers you are likely to see, with over **1,498 catalogued evidence files**.

In defamation law, what matters is not whether you click a particular button labelled "republish". What matters is that you:

- know about the content,
- know that it is alleged to be unlawful, and
- choose to keep hosting and displaying it to new readers.

Every time you send an "investigation complete - no breach" email, you are confirming publication with actual knowledge of the complaint. Your own wording that these are "one continuous moderation case" means each refusal to act is part of a single, ongoing decision to keep this review live.

That is the situation in which a platform ceases to be an innocent disseminator and becomes a full publisher in law.

5. Wider context: Trustpilot's role in an ongoing harassment campaign

As I have already told you:

- I am an individual Australian small-business owner with well under 10 employees.
- I maintain a structured harassment dossier of more than **1,498 evidence files** (emails, screenshots, PDFs, logs, payment records) covering the period **2019 to 2025**, documenting coordinated attacks across Discord, Facebook, Reddit and Trustpilot.
- Your platform has already:
 - removed and then quietly reinstated demonstrably false "scam" and "deletes your server" reviews after I provided technical and transactional proof,
 - received a formal cease and desist concerning specific defamatory content, and
 - been given examples of how competitors use your reviews to amplify fabricated narratives.

Your emails, internal ticket logs and moderation decisions in this case are being indexed directly into that chronology and will form part of any eventual regulatory or court record.

6. What I now require from Trustpilot

Given all of the above, I again require that you:

1. **Reopen case 00697534** and overturn the decisions that the "Hooded_saw" review is acceptable, including your prior findings under both "hate speech or discrimination" and "reviewing the wrong business".
2. **Remove the review in full**, on the basis that:
 - it is defamatory,
 - it constitutes disability-based harassment, and
 - it is not a genuine customer experience, but a retaliatory comment from a non-customer who breached my support rules and then left.
3. Record on my internal file that Trustpilot has been **explicitly told** it is already central to a potential multi platform defamation and (now) disability harassment action supported by a 1,498 file evidence base.
4. Confirm, in writing, what **concrete changes** you will make so that:

are no longer treated by Trustpilot as compliant "freedom of expression".

- reviews using mental health as an insult,
- in the context of documented harassment campaigns,
- and posted by non-customers,

If, after all of this, you still refuse to act, then I require you to confirm in writing that:

- you understand the review targets a named Australian small-business owner,
- you have been given detailed evidence that the wording is alleged to be unlawful defamation and disability-based harassment, and
- you are nonetheless choosing, as a deliberate policy decision, to keep publishing it.

That confirmation, together with this email chain and the attached evidence, will be preserved and may be relied upon in any future complaint to regulators or in any proceedings in which Trustpilot is joined as a publisher.

All rights are reserved.

Regards,

Anthony Brodie ("Soul")

Owner / Developer - Legion Killfeed

P.S. For completeness, I will also note that this is not happening in isolation. In August 2025 your own staff wrote to me stating that Trustpilot did **not** operate under an Australian Business Number, then within days reversed position and claimed to have an ABN and Australian entity all along. You have never properly explained that contradiction. In the same period you threatened to "defend and seek to recover costs" if I issued defamation proceedings, while quietly reinstating and then later removing the "WARZONE" review after previously insisting it would not be taken down – another clear example of republication after notice.

When I offered you a comprehensive evidence archive, you refused to access it, yet you continue to force my business to appear on your platform and demand internal proof that risks exposing customer information. That is why I have been careful about how much context I provide about who individual reviewers claim to be. You are compelling participation, publishing defamatory content, and at the same time asking for material that could compromise my customers' privacy, all against the backdrop of your own unresolved misstatements about your legal status in Australia.

To avoid any further suggestion that "format" or file-type limitations prevent proper review, I am providing the Discord chat logs both as a text-searchable PDF and as the original PNG screenshots, so your team can search the text directly and verify the timestamps for itself.

From: Soul Legion <legionkillfeed@outlook.com>
Sent: Thursday, 4 December 2025 11:26 PM
To: Trust Support <trust.operations@trustpilot.com>
Subject: Re: [thread::HUNjgNQfSqIVgYTI2NNAGA::]

Subject: Case 00697534 – Notice of continuing publication of disability-based defamation

Hi Megha,

Thank you for your email regarding case 00697534.

Your response again states that the review:

"Rude, unwilling to help, Soul needs professional help"

is neither defamatory nor discriminatory and that it aligns with Trustpilot's guidelines.

At this point you are not simply applying internal guidelines. You are choosing to continue publishing content that:

- targets me by name as a small business owner, and
- links my professional competence to an asserted mental illness ("Soul needs professional help").

1. Why this is disability-based harassment, not a normal "opinion" review

In ordinary English, telling someone they "need professional help" is widely understood as saying they are mentally ill or unstable and need psychiatric or psychological treatment. That is a statement about my mental health, not about your user's experience with my product.

Under the Disability Discrimination Act 1992 (Cth), mental illness is a form of disability, and harassment on the ground of disability in the context of goods, services and facilities is unlawful. The Act specifically makes harassment connected with the provision or use of services unlawful, including online services such as review platforms. ([Australian Human Rights Commission](#))

Your site is not simply recording neutral customer feedback here. It is hosting and republishing a personal attack about my supposed mental state, attached directly to my business listing.

2. Why the review is plainly defamatory

Australian defamation law uses a simple test: would the words tend to lower the person in the estimation of ordinary members of the community or cause others to avoid them. ([Judicial Commission of NSW](#))

Saying that a named sole trader "needs professional help" and is "rude" and "unwilling to help" clearly goes to my fitness to run my business. It suggests I am emotionally or mentally unfit to deal with customers and therefore unsafe or unsuitable to engage.

You have been shown the only communication that ever occurred with this person: the Discord direct messages in my previous email. Those messages contain nothing about mental health, diagnosis, treatment, or anything that could plausibly justify calling me mentally ill. The "professional help" allegation is not criticism based on disclosed true facts; it is an insult about my mental condition.

That means it cannot be justified as:

- a substantially true statement, or
- an honest opinion based on proper factual material,

because there is no factual basis at all for the suggestion that I have a mental illness that requires treatment. It is simply disability-based abuse tied to my business.

3. Why Trustpilot is now legally exposed as a publisher

In Australia, operators of pages and platforms that host third-party comments are treated as publishers for defamation purposes, even when someone else writes the words. The High Court has confirmed that opening a space for user comments can make you liable for any defamatory material in those comments. ([Gilbert + Tobin](#))

Platforms usually rely on the "innocent dissemination" defence. That defence requires that the intermediary:

- did not know the material was defamatory, and

- was not negligent in failing to remove it.

Australian commentary on that defence is clear: once a platform is put on notice that a review may be defamatory, it must act promptly or it risks losing that protection. ([Pentana Stanton Lawyers](#))

In my case, Trustpilot has now:

- received multiple flags and emails expressly identifying this review as both defamation and disability-based harassment;
- been given the full context of my interaction with the reviewer, showing there is no factual basis for the allegation about my mental health;
- repeatedly replied that the review is acceptable "freedom of expression" and that it will remain online; and
- issued "case closed" notices while continuing to display the review to new readers.

In other words, Trustpilot is now publishing this material with actual knowledge of the complaint and after being told that it is unlawful. That makes it very difficult for Trustpilot to argue that you lacked knowledge or acted reasonably if a court later finds that the review is defamatory and discriminatory.

4. Wider context and evidence

You are also aware that this review is not an isolated incident. I have already told you that:

- I am an individual Australian small-business owner, well under 20 employees.
- I maintain a structured harassment dossier of more than 1,400 evidentiary files across Discord, Facebook/Meta, Reddit and Trustpilot documenting a coordinated competitor-driven campaign.
- Your emails, ticket logs and moderation history around this specific review are being indexed into that chronology and will form part of any eventual regulatory complaint or civil proceedings.

Trustpilot is already at the centre of what is likely to become one of the most comprehensively documented multi-platform defamation and disability-harassment matters in Australia.

5. What I am now requiring

I again require that Trustpilot take one of the following steps:

1. Remove the review in full; or
2. At the absolute minimum, remove the words "Soul needs professional help" so that disability-based abuse is not being broadcast under your brand.

If Trustpilot refuses to do either, then I require you to confirm in writing that:

- you understand that the review targets a named Australian business owner;
- you have been expressly told that the wording is alleged to be unlawful defamation and disability-based harassment; and
- you are nonetheless choosing, as a deliberate policy decision, to keep publishing it.

That confirmation, together with this email chain, will be preserved as evidence and may be relied on in any future complaint to regulators or in any defamation proceedings in which Trustpilot is joined as a publisher.

All rights are reserved.

Regards,

ANTHONY BRODIE 
Owner/Developer - Legion Killfeed

From: Soul Legion <legionkillfeed@outlook.com>
Sent: Thursday, 4 December 2025 11:01 PM
To: Trust Support <trust.operations@trustpilot.com>
Subject: Re: [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Using DO NOT REPLY to avoid your legal responsibilities violates the law aswell, keep digging that hole, I'm coming with lawyers cunt.

From: Do not reply <donotreply@trustpilot.com>
Sent: Thursday, 4 December 2025 10:59 PM
To: legionkillfeed@outlook.com <legionkillfeed@outlook.com>
Subject: Update on your request #00697534 [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Hi,

Thank you for getting back to us. We would like to inform you that we have completed the assessment of your appeal.

After carefully reassessing the content, we found that the review is not defamatory. It reflects the reviewer's personal experience and opinion without making any unverified allegations of illegal or harmful conduct. As this falls within acceptable freedom of expression and does not violate our [Guidelines](#), the review will remain online.

We have concluded the investigation and we won't be responding further as we consider this case as closed .

Thank you for your understanding.

Kind Regards,

Megha



See our reviews


From: Soul Legion <legionkillfeed@outlook.com>
Sent: Thursday, 4 December 2025 10:36 PM
To: Trust Support <trust.operations@trustpilot.com>
Subject: Re: [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

I don't give a flying fuck about your external fucking evidence because this is the only communication that has taken place with this person, nothing about that implies I have a mental disability and you refusing to remove a post implying I need to seek professional help is unlawful. Now fucking remove it cunt.



Hooded

hooded.saw

This is the beginning of your direct message history with **Hooded**.

No servers in common

Add Friend

Block

Report Spam

November 29, 2025

1 ignored message — Hide



Hooded Sat 29 Nov 2025 7:27 AM

Good morning Soul,
I hope this message finds you well.
My name is Frank, writing from the Netherlands. I currently hold the admin position at Z Company. I understand you must be very busy operating such a large site, but I hope I can trouble you for a moment to clear the air.
You were recently in contact with my colleague, Lihou, and I'd like to provide a bit of background. Lihou is a 70-year-old British combat veteran who greatly enjoys playing DayZ it's his hobby. We're not experts in coding, but we are learning each day. At the moment, we are growing an active community and we rely on your support to maintain that growth.
After reading through the communication you had with my colleague today, I can see that things may have gotten off on the wrong foot. It seems we overlooked your request for clear information, and I fully understand that you can't perform your role properly without the necessary details.
With that said, I kindly ask that we continue the good service and cooperation we've enjoyed over the past months. To keep it brief: please let me know exactly what you require from us to get the feedback system back up and running.
If we are unable to work together on this, we may unfortunately have no choice but to consider moving our server funds to one of your competitors.
There will be no further communication regarding the points mentioned above from my end. The ball is in your court.
Kind regards,
Frank
Z Company

November 30, 2025



Soul  Legion Sun 30 Nov 2025 12:23 AM

I can't provide support to imbeciles who blatantly refuse to accept the help, using the internal ticket system and decided to throw a tantrum and have a scene.
This is the 5th time I've had this clown carry on and ignore tickets.

3 ignored messages — Hide



Hooded Sun 30 Nov 2025 2:36 AM

You are one sorry ass excuse
I'm laughing
Now block me

+ Message @Hooded



From: Trust Support <trust.operations@trustpilot.com>
Sent: Thursday, 4 December 2025 10:21 PM
To: legionkillfeed@outlook.com <legionkillfeed@outlook.com>
Subject: [thread::HUNjgfNQfSqIVgYTI2NNAGA::]

Case number: 00697534

Hi Anthony,

Thanks for your message. After further review of all communication regarding this issue, I can see that we've already come to a decision in an appeal previously raised in case **00697534**. I understand your frustration and I'm sorry that this is not the outcome you had hoped for or expected.

As part of the appeal process, your case was thoroughly investigated using all of the evidence that was available to us, or that we obtained from you or another party.

If you believe the review is against any other guidelines, you can flag it for the relevant reason. You can also refer to our Help Center article: ["For which reasons can businesses flag service reviews?"](#) This article provides further guidance on how and when to report reviews based on concerns.

In alignment with Trustpilot's guidelines, we arrived at the decision as stated in our previous message within the case **00697534**, and we won't be taking further action on this matter. This case is now closed.

Best regards,

From: Soul Legion <legionkillfeed@outlook.com>
Sent: Thursday, December 04, 2025 6:01 PM
To: Trust Support <trust.operations@trustpilot.com>
Subject: Re: Update on your request #00695808 [thread::SnAaDH8GdHu_wkCZLRsXNmA::]

For this single review, your own history now shows:

- 3 days ago: "Investigation complete: The review doesn't breach our guidelines for: hate speech or discrimination."

In other words, this is not a case where you are waiting for the right flag reason. You have already:

- rejected a flag that it is hate speech / discrimination (which is your own harassment category), and
- rejected a flag that it is even about the correct business, despite the text clearly targeting me personally by name ("Soul") and being attached to my listing.

On top of that, you have my emails explicitly stating that this is unlawful defamation and disability based harassment under Australian law. So in legal terms, you have already rejected it for harassment and defamation too. Your system history is now a clear record of repeated decisions to keep publishing it.

Whether you say "we never removed or republished it" is legally irrelevant. In defamation law, what matters is that you:

- know about the content,
- know that it has been alleged to be unlawful, and
- continue to host and display it to new readers.

Every time you conclude "investigation complete – no breach", you are confirming that publication with actual knowledge of the complaint. Your own email confirms that these are "one continuous moderation case", which means each refusal to act is part of a single, ongoing decision to keep it live.

You already know:

- I am an individual Australian small business owner (well under 10 employees).
- The words "Rude, unwilling to help" and "Soul needs profession help" are attached to my Trustpilot profile and presented as a trusted review to potential customers.
- They carry imputations that I am professionally incompetent and that I am mentally unwell or unstable and therefore unfit to operate my business.

Under the uniform Defamation Acts, that is classic defamation of an individual and an excluded corporation. Under the Disability Discrimination Act 1992 (Cth), using an imputed mental illness as an insult in the course of providing services to an Australian user is disability based harassment, not "robust opinion".

You also know this is not an isolated review. I have already told you that:

- I maintain a structured harassment dossier of more than 1,498 evidence files (emails, screenshots, PDFs, logs, payment records) across Discord, Facebook/Meta, Reddit and Trustpilot, covering 2019 to 2025.
- That dossier shows a coordinated, competitor driven campaign, with Trustpilot acting as a central amplification point through review bombing, quiet reinstatements, and refusal to act on clearly abusive content once notified.
- Your emails, ticket logs and the moderation history for this review are already indexed in that chronology and will form part of any eventual proceedings.

So you are not dealing with a simple guideline query. You are dealing with a review you have:

- been on notice about since the first flag,
- re assessed multiple times under different guideline categories, and
- nonetheless chosen to keep publishing, despite being told in plain language that it is unlawful defamation and disability based harassment and that it sits inside one of the most comprehensively documented multi platform defamation cases you are likely to see.

At this point, the position is straightforward:

- Your internal guideline labels do not override Australian law.
- Your "we only assess for the reason selected" rule does not protect you once you are on notice that the content is alleged to be illegal. You now have more than enough information to act without any further drop down selections from me.

Again, I require that you:

1. Reopen and overturn the decisions that the Hooded_saw review is acceptable, including your findings under both "hate speech or discrimination" and "being about a different business".
2. Remove the review in full on the basis that it is defamatory and constitutes disability based harassment, not merely as a discretionary removal "at the business's request".
3. Record on my internal file that Trustpilot has been expressly told it is already central to a potential multi platform defamation action supported by a 1,498 file evidence base.
4. Confirm, in writing, what concrete changes you will make so that reviews using mental health as an insult, in the context of a documented harassment campaign, are no longer treated by Trustpilot as compliant content.

You now cannot credibly say this is a misunderstanding about flag types. If you continue to publish this review in its current form, it will be treated as a deliberate decision to keep clearly abusive and unlawful material online with full knowledge of the consequences.

Regards,
Anthony Brodie (Soul)
Legion Killfeed

From: Soul Legion <legionkillfeed@outlook.com>
Sent: Thursday, 4 December 2025 6:01 PM
To: Trust Support <trust.operations@trustpilot.com>
Subject: Re: Update on your request #00695808 [thread::SnAaDH8GdHu_wkCZLRsXNmA::]

For this single review, your own history now shows:

- 3 days ago: "Investigation complete: The review doesn't breach our guidelines for: hate speech or discrimination."

In other words, this is not a case where you are waiting for the right flag reason. You have already:

- rejected a flag that it is hate speech / discrimination (which is your own harassment category), and
- rejected a flag that it is even about the correct business, despite the text clearly targeting me personally by name ("Soul") and being attached to my listing.

On top of that, you have my emails explicitly stating that this is unlawful defamation and disability based harassment under Australian law. So in legal terms, you have already rejected it for harassment and defamation too. Your system history is now a clear record of repeated decisions to keep publishing it.

Whether you say "we never removed or republished it" is legally irrelevant. In defamation law, what matters is that you:

- know about the content,

- know that it has been alleged to be unlawful, and
- continue to host and display it to new readers.

Every time you conclude "investigation complete – no breach", you are confirming that publication with actual knowledge of the complaint. Your own email confirms that these are "one continuous moderation case", which means each refusal to act is part of a single, ongoing decision to keep it live.

You already know:

- I am an individual Australian small business owner (well under 10 employees).
- The words "Rude, unwilling to help" and "Soul needs profession help" are attached to my Trustpilot profile and presented as a trusted review to potential customers.
- They carry imputations that I am professionally incompetent and that I am mentally unwell or unstable and therefore unfit to operate my business.

Under the uniform Defamation Acts, that is classic defamation of an individual and an excluded corporation. Under the Disability Discrimination Act 1992 (Cth), using an imputed mental illness as an insult in the course of providing services to an Australian user is disability based harassment, not "robust opinion".

You also know this is not an isolated review. I have already told you that:

- I maintain a structured harassment dossier of more than 1,498 evidence files (emails, screenshots, PDFs, logs, payment records) across Discord, Facebook/Meta, Reddit and Trustpilot, covering 2019 to 2025.
- That dossier shows a coordinated, competitor driven campaign, with Trustpilot acting as a central amplification point through review bombing, quiet reinstatements, and refusal to act on clearly abusive content once notified.
- Your emails, ticket logs and the moderation history for this review are already indexed in that chronology and will form part of any eventual proceedings.

So you are not dealing with a simple guideline query. You are dealing with a review you have:

