

From: Soul Legion <Brodiess@REDACTED.email>
To: Trust Support <trust.operations@trustpilot.com>, Trustpilot Legal <legal@trustpilot.com>, Trustpilot Content Integrity <contentintegrity@trustpilot.com>, Press - Trust Pilot <press@trustpilot.com>
Subject: Re: Update on the review you flagged - Case #00709812
Date: Sat, 06 Dec 2025 12:18:03 +1100

Subject: Follow-up: non-customer reviews, flagging suspension, and privacy / safety concerns (Kamy & "Hooded_saw")

HiTrustpilot team,

This is a further follow-up to my earlier escalation regarding the reviews on my Legion Killfeed profile and your email of 5 December about Case #00709906. You now appear to have removed the "Unholy" review, which underlines that my concerns were well-founded. At the very same time, you have suspended my ability to flag reviews for 48 hours, while leaving two remaining reviews online from:

- **Kamy** – a direct competitor, not a customer; and
- **Consumer / "Hooded_saw"** – an unverified third party who has never been my customer and whose only interaction with me was via rule-breaking DMs.

I want to make it absolutely clear what that combination now looks like in privacy, safety, and legal terms.

1. "Hooded_saw" is still a non-customer, and your own evidence shows it

The "Consumer / Hooded_saw" review is still published as if it were a genuine customer experience. It is not.

- The only paying subscriber in this matter is the older server owner using the Discord account @lihou2451.
- The reviewer's Apple relay address 2znkkydr98@privaterelay.appleid.com does **not** appear anywhere in my Ko-fi membership records, for any tier, including ended subscriptions.
- The handle "**hooded_saw**" also does **not** appear in my Ko-fi membership list.
- At no time has this person opened a ticket, authenticated in my system as a customer, or been confirmed by the actual subscriber as an authorised contact.

What you actually have is an unverified stranger, using a proxied email address, trying to bypass my support rules in private DMs to gain access to a paid server that belongs to someone else. When I refuse to break my own privacy and security obligations and insist that support must occur in the ticket the subscriber already has, that stranger runs to Trustpilot and is allowed to publish a "review" attacking my conduct. From a privacy and security perspective, that is reckless. I am not allowed to treat a random person with a relay email as if they were my customer, especially when elderly subscribers are involved. Yet you are treating that same person as a "verified user" with the right to publicly rate my business.

If "Hooded_saw" truly is connected to the elderly subscriber, then your platform is helping them exert pressure **outside** the proper support channel and potentially exploit that subscriber's account and payment details. If they are **not** connected, you are hosting a non-customer's attack as if it were a genuine service experience. In either case, it is not a lawful or honest "consumer review".

2. Elder-exploitation risk is not theoretical

You have already been told that the only subscriber in this scenario is described by his friend as a 70-year-old combat veteran. You also know that:

- He ignored his own ticket instead of answering a basic support question.
- His "admin" friend then tried to obtain access through DMs.
- That friend is now the one writing the Trustpilot review, while still refusing to be accountable in the actual ticket.

That pattern is exactly what raises elder-exploitation concerns. A technically capable "friend" interposes themselves, pressures a service provider to break normal security practice, and then weaponises an older person's subscription when those boundaries are enforced.

I have already indicated that I will be contacting the eSafety Commissioner and other appropriate authorities again, and this situation will be included. Your decision to side with an unverified DM-only "friend" over the basic privacy and security of an elderly subscriber is going straight into that record, alongside the rest of your correspondence and the 400+ emails that now document your handling of this saga, including roughly **380 emails in a single month** during the August–September escalation.

3. Your "AI-powered defamation tool" is clearly not assessing context

You repeatedly tell me that reviews have been "run through" an "AI-powered defamation tool" and that it has decided they are not defamatory. That claim cannot be reconciled with what is actually happening.

A genuine AI system, operating in good faith, would have:

- Recognised that **Kamy** is a direct competitor in the same niche, with a long documented history of harassment, and treated his "Another bad review" as a conflicted, non-consumer attack.
- Picked up the striking repetition in language and themes across the three clustered reviews (Kamy, Unholy, and Consumer/Hooded_saw), including that two were posted **after** I flagged the original and clearly function as amplification and retaliation rather than independent feedback.
- Correlated these reviews with the extensive evidence bundle you already hold showing the wider smear campaign and the commercial rivalry involved.
- Flagged the disability-based insult and personal character smears as high-risk content instead of telling me that "Soul needs professional help" somehow "doesn't breach our guidelines".

Instead, your pattern is consistent with a very simple keyword filter that only reacts if certain magic words appear, and otherwise defaults to "not defamatory" no matter how much context is supplied. There is no sign that your system is genuinely evaluating harm, repetition, motive, or commercial conflict. At best, that is grossly negligent. At worst, your "AI" branding is a marketing veneer over a tool deliberately tuned to leave as much damaging content online as possible regardless of the law.

4. You removed Unholy but then blocked my only internal remedy

The removal of the **Unholy** review confirms that your earlier "no breach" assessment was wrong. It was defamatory and harmful enough to warrant deletion once I pushed hard enough and supplied more evidence.

Immediately after that, you:

- suspended my ability to flag reviews for 48 hours, citing "misuse of the flagging tool"; while

- leaving the competitor review (**Kamy**) and the non-customer review (**Consumer/Hooded_saw**) online; and
- continuing to urge me to simply “tell my side of the story” in replies.

You are therefore:

- actively restricting my ability to use your own safety mechanism;
- doing so in direct response to my attempts to challenge clearly abusive and commercially conflicted content; and
- placing the burden entirely on me to “respond” publicly while your platform presents these reviews as neutral, vetted customer experiences.

That is not a neutral moderation decision. It is an aggravated harm: you have become a bottleneck that both amplifies defamatory reviews **and** blocks the victim’s only internal mechanism for challenging them.

For the avoidance of doubt, my earlier subject line about killing myself because of Trustpilot’s abuse was not a plan to self-harm. It was an attempt to convey how extreme and unsustainable the pressure from your mishandling has become. I do not want to die; I want the harassment and your role in enabling it to stop. Your current course of conduct is pushing me into suicidal ideation, and if this is what it does to a 32-year-old adult with extensive documentation skills, I am deeply concerned about what it could do to less resilient people and to children.

5. What I now require

Given all of the above, my position is:

Immediate removal of the remaining two reviews (Kamy and Consumer/Hooded_saw) in full, on the combined basis that:

- Kamy is a direct competitor, not a customer, and is using your platform to tell potential customers to choose rival bots;
- Consumer/Hooded_saw is not, and has never been, a verified customer, and his review relies on a false claim of being a “paying customer” to attack me for enforcing basic privacy and support rules.

1. **Restoration of my flagging ability** on my own profile, and confirmation that you will not penalise me for flagging coordinated harassment and non-customer reviews where there is substantial evidence of abuse and commercial conflict.

2. A clear written explanation of how your so-called “AI-powered defamation tool” actually operates in cases like this, including:

At present, its behaviour is inconsistent with any genuine, context-sensitive AI system, and I consider that discrepancy relevant to both regulators and any future legal proceedings.

- whether it looks at cross-review patterns and prior evidence you hold, or
- whether it simply runs per-review keyword checks.

3. Confirmation that you understand and accept that allowing an unverified DM-only “friend” to publicly rate my service as if they were the paying customer, in circumstances where an elderly subscriber’s account may be being misused, raises serious privacy, consumer-protection, and duty-of-care issues for Trustpilot, not just for me.

This email, your response, and the existing body of more than 400 emails between myself and Trustpilot (including the ~380 exchanged during the August–September blow-up) will be filed with the eSafety Commissioner and retained for any future legal or regulatory process examining your conduct.

Regards,

Anthony Brodie (Soul)

Owner and developer, Legion Killfeed



From: Soul Legion <Brodies@REDACTED.email>

Sent: Saturday, 6 December 2025 6:30 AM

To: Trust Support <trust.operations@trustpilot.com>; Trustpilot Legal <legal@trustpilot.com>; Trustpilot Content Integrity <contentintegrity@trustpilot.com>; Press - Trust Pilot <press@trustpilot.com>

Cc: Trustpilot Content Integrity <contentintegrity@trustpilot.com>; Trustpilot Legal <legal@trustpilot.com>; Press - Trust Pilot <press@trustpilot.com>

Subject: Re: Update on the review you flagged - Case #00709812

You have 45 minutes, then I make a noose, you want me to play chicken with these cunts, chicken it is with you, because I will do it. I've had enough.

Play checked now cunt.



From: Soul Legion <Brodies@REDACTED.email>

Sent: Saturday, 6 December 2025 6:27 AM

To: Trust Support <trust.operations@trustpilot.com>; Trustpilot Legal <legal@trustpilot.com>; Trustpilot Content Integrity <contentintegrity@trustpilot.com>; Press - Trust Pilot <press@trustpilot.com>

Subject: Re: Update on the review you flagged - Case #00709812

Subject: Case #00709906 – Retaliatory suspension of flagging and continued publication of coordinated competitor abuse - REMOVE IT OR I SUICIDE.

<https://www.youtube.com/embed/rcMqoRAILpk?si=gAVikA1-q681Jq>

Hi Gowtham,

You’ve now moved beyond bad moderation and into outright retaliation.

Within the space of a few hours you have:

- Confirmed that you will keep publishing a defamatory review from a direct competitor, “Kamy”, while pretending he is a normal customer.
- Left standing the “Unholy” review, which is another hit piece from the same circle, telling people to use other bots.

- Continued to host the “Consumer / Hooded_saw” review from a non-customer whose only interaction with me was rule-breaking DMs and disability-based abuse.
- Suspended my ability to flag reviews for 48 hours for supposed “misuse of the flagging tool”.

Let’s be clear on what actually happened.

I flagged:

- Kamy’s review – a competitor in the same niche, using your site to trash my business and recommend “elsewhere”.
- Unholy’s review – a character-assassination rant explicitly steering people to other bots.
- Consumer/Hooded_saw – an unverified “friend” of an elderly subscriber who never opened his own ticket, never paid for anything, never used the service through any channel I can verify, and whose “experience” is limited to harassing me in DMs after the real account holder ignored an open ticket.

That is not “misuse of the flagging tool”. That is exactly what the tool exists for.

Your own pop-up claims I was “flagging at a high speed”. In reality I was flagging a small cluster of obviously connected reviews, all tied to the same people you have already been told about, in writing, for months. The “high speed” you are punishing is nothing more than me responding promptly to a coordinated attack on my business.

Suspending my flagging rights while you continue to publish those same reviews has three practical effects:

It obstructs my ability to put you on notice about new defamatory or abusive content.

1. It allows you to say later that “we weren’t notified” while you are actively throttling the only structured reporting channel businesses have.
2. It sends a very clear message that if a small business uses your systems to push back on coordinated harassment, you will simply turn the tools off and keep the harassment live.

That is not a good-faith application of guidelines. It is cover for continuing publication after notice.

Your behaviour needs to be looked at in context, which you keep trying to ignore:

- You told me in writing earlier this year that Trustpilot did not have an Australian Business Number and didn’t need one.
- After I pointed out the legal problems with that, an ABN suddenly appeared in your disclosures with no explanation and no correction of the earlier statements.
- You have a long history on my account of removing obviously false reviews, then quietly reinstating them later.
- In the Consumer/Hooded_saw case, you insisted that “Soul needs professional help” was compliant, then allowed the reviewer to silently edit that disability-based insult out rather than admitting it should never have been published. That is republication after notice, no matter how you spin it.
- Now, when the same small group starts up again under new names, you not only keep their reviews live, you also punish me for using your own reporting tools.

This is exactly why I do not trust your “AI-powered defamation tool”, your internal guideline labels, or your assurances that these reviews are harmless “opinions”. Your track record on my business – including your false statements about your ABN status – shows that your classifications are driven by internal convenience, not by any serious attempt to comply with Australian law or protect targets of harassment.

For the avoidance of doubt, my position is:

- Kamy is a competitor, not a customer. Leaving his review up is knowingly hosting a false “consumer” review from a commercial rival.
- Unholy is clearly part of the same harassment network, using your platform to repeat the same narratives and drive people to other bots.
- Consumer/Hooded_saw is not and has never been my customer and was never an authorised contact for the 70-year-old subscriber he claims to speak for. You have the full ticket history and DM logs showing that.
- Your 48-hour flagging suspension is retaliatory and designed to shut down further notice while you continue publishing material you already know is disputed and potentially unlawful.

At this point I am done playing along with the idea that this is just “guideline disagreements”.

Unless you:

- remove the Kamy, Unholy and Consumer/Hooded_saw reviews in full as non-genuine, malicious reviews; and
- lift the suspension on my flagging tools and record that they were being used for their intended purpose,

I will proceed on the basis that:

- Trustpilot has chosen to side with a coordinated competitor-driven harassment campaign against a small Australian business;
- you are continuing to republish defamatory material after repeated notice, while actively obstructing further reporting; and
- your earlier ABN misrepresentations and your contradictory handling of prior reviews will be part of the evidence of systemic bad faith.

Any further responses that simply restate “our guidelines say it’s fine” or try to blame an “AI tool” will be treated as more of the same and added to the litigation and regulatory file. I am not wasting any more time arguing with a system that punishes me for using the tools you told me to rely on.

Regards,

Anthony Brodie (Soul)

Owner and developer – Legion Killfeed

<https://www.youtube.com/embed/rcMqoRAILpk?si=gAVikA1-q681Jq>



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

Sent: Saturday, 6 December 2025 4:03 AM

To: Brodies@REDACTED.email <Brodies@REDACTED.email>

Subject: Update on the review you flagged - Case #00709812

Hi killfeed.co,

We’re following up about the [review](#) by Consumer, which you flagged as not based on a genuine experience.

Our decision

We've carefully assessed the review based on your selected flagging reason. We've found that the review follows our guidelines, so it will remain online.

How we assess flagged reviews

When reviews are flagged for not being based on a genuine experience, we run them through [a series of checks](#) and analyze various factors, such as:

- The reviewer's activity on Trustpilot (e.g., number of previous reviews and past engagement with businesses)
- If the content of the review follows our guidelines for reviewers
- Patterns of unusual behaviour detected by our automated technology

In this case, we weren't able to detect any suspicious patterns or content, so the review stays online.

What counts as a genuine experience

[A genuine experience](#) doesn't always mean a purchase. It could also include an interaction, like:

- Browsing your website
- An online chat with your business

Reviewers don't need to use their real names when submitting their review, so you might not always recognize them from the reviewer name alone. This is why we recommend using our "[request information](#)" tool or replying to the review so you can share your perspective.

Share your side of the story

If you haven't already, we recommend [replying to the review](#). Your reply appears alongside the review and gives potential customers a complete picture from both your side and the reviewer's—which can help build trust and credibility.

Want us to take another look?

If you believe we've made a mistake, please reply to this email with any additional context, and our team will reassess the case.

Thanks,

Trustpilot

Please note that depending on the situation, we may use automated detection technology, a team of experts, or a combination of both to assess content on the Trustpilot platform.

Email Headers (Raw)

X-Mozilla-Status	0001
X-Mozilla-Status2	00000000
From	Soul Legion <Brodies@REDACTED.email>
To	Trust Support <trust.operations@trustpilot.com>, Trustpilot Legal <legal@trustpilot.com>, Trustpilot Content Integrity & It;contentintegrity@trustpilot.com>, Press - Trust Pilot <press@trustpilot.com>
Subject	Re: Update on the review you flagged - Case #00709812
Thread-Topic	Update on the review you flagged - Case #00709812
Thread-Index	AQHcZgkcBT32eXhvJEqawgwUMZc5RrUTZ4blgAAG9SWAAEUOMg==
X-MS-Exchange-MessageSentRepresentingType	1
Date	Sat, 06 Dec 2025 12:18:03 +1100
Message-ID	<PU4P216MB15689C9D9FC63D797FF7A2C1DAA7A@PU4P216MB1568.KORP216.PROD.OUTLOOK.COM>
References	<emUxk00T6T3E0004syodOphSC-4UKdakmxQLg@sfdc.net> <PU4P216MB1568FCD60EF39382C6FFF6B7DAA7A@PU4P216MB1568.KORP216.PROD.OUTLOOK.COM> <PU4P216MB1568CB4AF8CAB70459C3DD3EDAA7A@PU4P216MB1568.KORP216.PROD.OUTLOOK.COM>
Keywords	Trustpilot
In-Reply-To	<PU4P216MB1568CB4AF8CAB70459C3DD3EDAA7A@PU4P216MB1568.KORP216.PROD.OUTLOOK.COM>
Content-Language	en-AU
X-MS-Has-Attach	yes
X-MS-Exchange-Organization-SCL	-1
X-MS-TNEF-Correlator	
X-MS-Exchange-Organization-RecordReviewCfmType	0
msip_labels	
x-ms-exchange-imapappendstamp	PU4P216MB1568.KORP216.PROD.OUTLOOK.COM (15.20.9388.012)
Content-Type	multipart/related; boundary="_019_PU4P216MB15689C9D9FC63D797FF7A2C1DAA7APU4P216MB1568KORP_"; type="text/html"
MIME-Version	1.0