




Reverse engineering hate

Liz Fong-Jones, on behalf of End Kiwi Farms
2023-11-07



What is Kiwi Farms?

A web forum that facilitates the discussion and harassment of online figures and communities. -



Targets are often subject to organized group trolling and stalking, as well as doxxing and real-life harassment. -



Hundreds

of threads on individuals

3

documented suicides of
victims (all trans/NB)

Typical consequences of being targeted

Personal information doxxed (including government ID numbers, home address)

Family members, friends, acquaintances, colleagues also named and targeted

Massive copyright infringement of photographs & defacement for derogatory purposes

“Fanfiction” written to create and refine lies to cast individual in worst possible light

Unwanted contact via email, telephone, and sometimes in-person, threatening harm

Unwanted contact spreading lies to employer, landlord to cause firing & eviction

Attempts to poison SEO to make unhireable

Thousands

of collateral targets

“When something is ours, it’s ours forever. We might get bored with it and put it down. We might play with it a little too hard and break it. But it never stops being ours, and someday in the future we’ll remember it, and pick it back up again. Would you like to become ours, my dear?”

How is Kiwi Farms hosted?

How most sites are hosted

All-in-one package - domain, servers in a datacentre, IP addresses, network transit, DNS servers, email MX, etc

Optional DDoS protection services by a reputable provider eg Fastly, Cloudflare, Akamai, GCP/Amazon etc

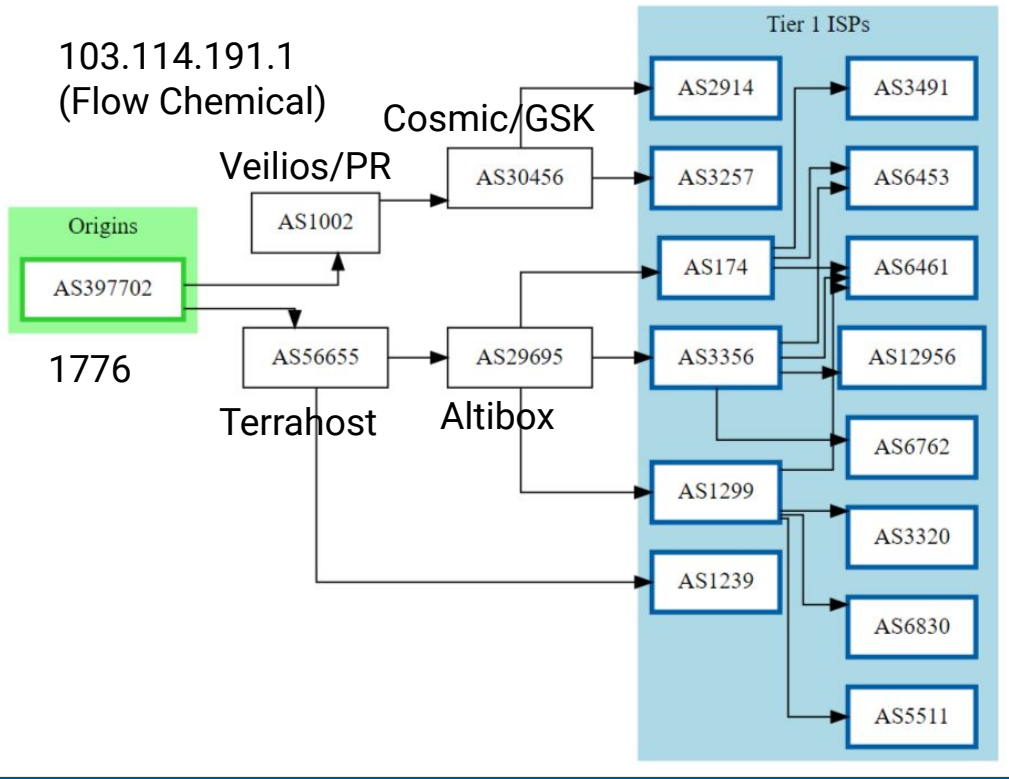
At most, 3-5 different providers in use, not dozens.

Hate sites migrate to shithosts/self-hosting

Each component obtained a la carte from cooperating or unwitting hosts (often “bulletproof” hosters)

Shell companies for IP addresses, ASNs/BGP routing, network contracts for transit, subleasing datacentre floor space, hand-written DDoS protection, different providers for frontend vs backend vs MX

Payment ban evasion/laundering; multiple nested LLCs



Domain reg by Epik

DNS servers on VSYS
(AS30860) & AS208046

Additional peering via
AS30860 & AS208046

MX with Terrahost

Physical servers w/
Terrahost & VSYS

Volumetric defence by
Terrahost & Cosmic

Home-grown PoW
app-layer defence

Example network diagram circa November 2022

Extreme hate speech
& malware are alike.

It looks like a duck,
quacks like a duck, yet
isn't recognised as one.

Key problems in obtaining action

Plausible deniability of
off-forum coordination

“In Minecraft” & other slang
to avoid casual scrutiny

Sensitive info hidden behind
registration gates

“Freedom of speech”
“Section 230”

“We don’t send unsolicited
IP traffic”

“Net neutrality”

“There’s worse stuff
elsewhere”

How does deplatforming
work?

First, on the ethics

Hurricane Electric Acceptable Use Policy (AUP)

All Hurricane Electric (Hurricane) colocation, dedicated server, or direct internet connection customers are bound by the following Acceptable Use Policy. This document may be updated from time to time. Please consult this site periodically for the most recent revision of this document.



products & services solutions network about cogent support offices

Sign In

CONTACT US

product or service that

n reasonably
ul, or excessively
te human violence

y, publicity, or other
ng patents,
that would infringe,
ices, it may not be

CONTACT

nful to or interfere
ntent includes, but is
ure programs.
ing fraudulent offers
ding statements,
on any order form,

es, including, without
; and (ii) send e-mail
hibited from using
n addition, Users will
does not wish to
cluding, without

or receive e-mail messages in a manner that violates the use policies of any other Internet
box exclusively as a storage space for data.

Third party links, content. Users will not use the Services in violation of the rules, guidelines or agreements associated with
search engines, subscription Web services, chat areas, bulletin boards, Web pages, USENET, or other services accessed via the
Services.

Inappropriate Actions. Users will not use the Services to conduct activities that may be harmful to or interfere with the
Services, a User's terminal session or any third party's networks, systems, services, or websites. Users will not engage in any
activities designed to harass, or that will preclude or interfere with the use of Service (e.g., synchronized number sequence
attacks) by any other User on the Lumen network or on another provider's network. In addition, Users will not use the Service
(a) by any means or device to avoid payment; (b) to access User's account or Lumen Services after User has terminated User's
account; (c) on behalf of persons or firms listed in the Spamhaus Register of Known Spam Operations database at

not to engage in any unacceptable use of the service.

What Uses are Prohibited?

Unacceptable use includes, but is not limited to, any of the following:

1. Posting, transmission, re-transmission, or storing material on or through any of Cogent's products or services, if in the sole judgment of Cogent such posting, transmission, re-transmission or storage is: (a) in violation of any local, state, federal, or non-United States law or regulation (including rights protected by copyright, trade secret, patent or other intellectual property or similar laws or regulations); (b) threatening or abusive; (c) obscene; (d) indecent; or (e) defamatory. Each customer shall be responsible for determining what laws or regulations are applicable to his or her use of the products and services.

2. DENIALS - NO CUSTOMER

Hurricane's IP space (or Customer) put on black by Spamhaus (<http://www.spamhaus.org/>) or other similar organizations, or perform activities that would cause portions of the Internet to block mail or refuse to route traffic to any portion of Hurricane's IP space (or address space announced by Hurricane on behalf of Customer).

3. Perform actions that cause unusual load on Hurricane servers (for example, mail servers, web servers, usenet servers, name servers, etc.), that cause slowness or denial of service to other Hurricane customers.

3. Do anything that threatens the Internet or any other network.

1. No customer shall take actions that cause any portion of the Internet, or the Internet as a whole, to become unusable to any other portion of the Internet, or the Internet as a whole.

for any violation of the Policy. Zayo also reserves the right to avail itself of the safe harbor

or receive e-mail messages in a manner that violates the use policies of any other Internet
box exclusively as a storage space for data.

Initially: Tweet storms.

(but this stopped working)

The image shows a vertical stack of three tweets from the account @voxility, dated October 25, 2022. Each tweet is a reply to @lizthegrey. The first tweet has 1 reply, 18 likes, and 1 retweet. The second tweet has 15 replies, 13 retweets, and 50 likes. The third tweet has 18 replies, 13 retweets, and 75 likes. The tweets describe the blocking and subsequent permanent banning of the entity kiwifarms.net from the Voxility network.

Voxility @voxility · Oct 25, 2022
Replying to @lizthegrey
[@lizthegrey](#) thank you for letting us know. Our networking & abuse teams are looking at this.

Voxility @voxility · Oct 25, 2022
Replying to @lizthegrey
Corresponding IP for [kiwifarms.net](#) - transiting Voxility network - has been blocked. We will continue to monitor this entity. Thank you!

Voxility @voxility · Oct 25, 2022
Replying to @lizthegrey
They will most likely find other means to come back online. Entity [kiwifarms.net](#) is now permanently banned across Voxility Network and there are multiple alerts in place to notify our abuse team if they resurface using the same subnet or different hosting reseller.

Complain to origin

Yes, they will retaliate.

At best, with publishing the complaint.

At worst, with doubling down.


Interior Take that off the god damn Internet!

2017-08-18 - Liz Fong-Jones: Harassing people for hosting my email.

Thread Status: Not open for further replies.

1 2 3 4 5 6 - 8 Next >

The Lolcow Wiki has an article on Trans Lifeline

 Null

kiwifarms.net

Liz Fong-Jones is some tranny with a jaw sharper than Robert Z'Dar who Greta Gustava / Kjel Andeson has apparently employed to harass third parties vaguely associated with me. This is to, presumably, help cover up that Trans LifeLine is a scam. Zhe works for Google, and is using an authoritative @google.com address to do this network for Kjel. Xar's @google.com addresses are good for scaring small business.

This email falsely accuses me, us, whomever, of criminal behavior. No one is blackmailing this fucking troll. The only thing I'd want them to do is fuck off. I don't know if they're intimidated and I don't care. They should be more worried of the fact that they are openly using their @google.com association for non-work related activities, such as fucking with people they don't know and who never heard of them until today.

Page 1 / 10

Then wa
for respo
non-resp

#908447 - We Have Answered Your Ticket - Terrahost



Terrahost AS <noreply@terrahost.no>

To: You

Hello,

This is a notification to let you know that one of our support representatives has answered your ticket.

The message is:

Hello ~~_____~~

Thank you for notifying us. We already know, and we do not care. Thank you!

Regards,
TH Support

To reply to this ticket, please log into ENIGMA.

Thank you for using Terrahost.

Go further up the chain

Use evidence of previous refusals down the chain.

Cite specific evidence & AUP clauses violated.

Wait longer and longer at each step, as providers are busier/more indifferent.

Abusive netblock 103.114.191.0/24, kiwifarms.net, from AS397702 propagated by AS29695
1 message

Liz Fong-Jones <lizf@honeycomb.io> Fri, Nov 4, 2022 at 9:15 PM
To: abuse@altibox.no

To whom it may concern,

I am con
complain
I've also **Abusive netblock 103.114.191.0/24 (kiwifarms.net) from AS397702 propagated by AS174**
2 messages

In brief, I
scrambli
technolo **Liz Fong-Jones** <lizf@honeycomb.io> Fri, Nov 4, 2022 at 8:22 PM
To: abuse@cogentco.com

To whom it may concern,

They are
use of th
directly f I am contacting you regarding abusive content originating from AS397702 transiting through Cogent via Terrahost (AS56655) & its affiliate Altibox (AS29695). Terrahost, a subsidiary of Epik, is unresponsive to abuse complaints (and in fact exposes to further doxing anyone who might file complaints). Altibox similarly has willingly taken on propagation of these routes knowing exactly the type of content served from them.

As soon
access li
forwarde In brief, Kiwi Farms is a site dedicated to stalking and harassing vulnerable people, especially transgender and neurodivergent people. They've been in the news recently, and are scrambling since they lost Cloudflare's protection: <https://www.nbcnews.com/tech/internet/cloudflare-kiwi-farms-keffals-anti-trans-rna44834> <https://www.washingtonpost.com/technology/2022/09/03/cloudflare-drops-kiwifarms/> among other coverage.

The net
(1776 St
upstrear They are the epitome of a violation of Cogent's AUP forbidding "1. (b) threatening or abusive; (c) obscene; (d) indecent; or (e) defamatory [content] and 15. harassment, whether through language, frequency, or size of messages", but thus far they've escaped scrutiny by being a customer of a customer of Cogent directly paying Cogent for bandwidth.

It appea
(the AS)
Terrahos
further s As soon as Zayo discovered who they had signed a direct contract with, Kiwi Farms/1776 Solutions were terminated for ToS violations three weekends ago. Voxility likewise [blocked access](#) last week for AUP violation. GTT Communications recently blocked forwarding of route announcements for Kiwi Farms as of earlier this week. And Path Networks haven't forwarded route announcements for 1776 Solutions since last year: https://twitter.com/path_network/status/1566322558672244742

#90844
TA The netblocks and IP addresses at issue are [103.114.191.0/24](#), specifically [103.114.191.1](#) which Kiwifarms.net currently resolves to. The ASNs currently involved are AS397702 (1776 Solutions, run by Joshua Moon, the proprietor of Kiwi Farms), AS56655 (Terrahost, a subsidiary of Epik), and AS29695 (Altibox, an affiliate of Terrahost). In this fashion, Joshua Moon evades any direct enforcement of upstream AUPs while still being able to use your network.

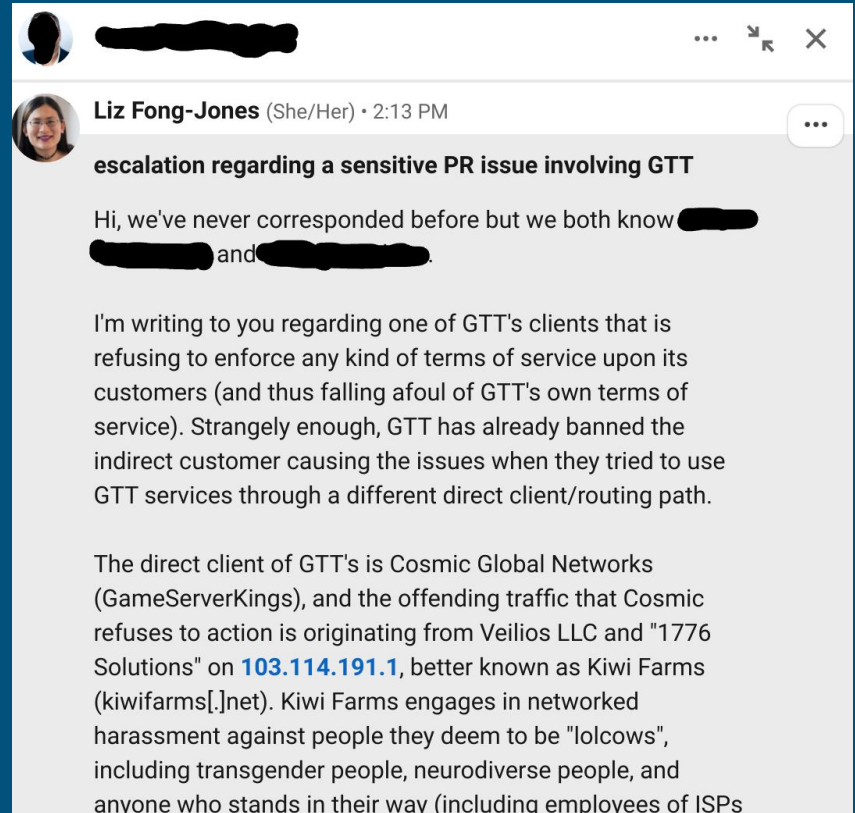
#908447 - We Have Answered Your Ticket - Terrahost

TA Terrahost AS <noreply@terrahost.no>
To: You
Hello,
This is a notification to let you know that one of our support representatives has answered your ticket.
The message is:

Wait multiple
weeks per step in
chain if no action.

Find escalation points

Email/LinkedIn & network
to find individuals who
will use internal
escalation paths.



If all else
fails, litigate.



[Print subscriptions](#) [My account](#) [Search jobs](#) [Search](#) [Australia edition](#)

Support the Guardian
Fund independent journalism with \$5 per month
[Support us](#) →

The Guardian
A decade of making a difference

[News](#) [Opinion](#) [Sport](#) [Culture](#) [Lifestyle](#) [More](#)

[UK](#) [World](#) [Climate crisis](#) [Ukraine](#) [Football](#) [Newsletters](#) [Business](#) [Environment](#) [UK politics](#) [Education](#) **Society**

Online abuse
Defamation in the internet age: could a \$400,000 Australian court ruling silence the notorious online forum Kiwi Farms?

An Australian company that helped the website remain accessible has been ordered to pay damages - but the legal arguments remain unsettled

- [Follow our Australia news live blog for latest updates](#)
- [Get our morning and afternoon news emails, free app or](#)

This process is broken.



Our team has had more success, at a cost

12 volunteers

16 months of work &
counting

42 providers dropped

1000s of hours spent

4 parallel legal teams in
US, CA, AU, IE

Approx 100,000 GBP
spent on litigation,
excluding pro bono
assistance

Feedback on “No Easy Way Out”

Publicly observable effort != private efforts

Confounding factors (e.g. “new” users may be those that lost passwords and had to recreate)

Measuring a months/years-long process after 4 months is premature.

Talk to malware/cybercrime takedown experts.

Considerations on funding/biz model and whether KF/Moon are an exceptional case

Comparison of “cow” vs “social” subforums?

No Easy Way Out: the Effectiveness of Deplatforming an Extremist Forum to Suppress Hate and Harassment

Anh V. Vu
University of Cambridge
Cambridge Cybercrime Centre
anh.vu@cl.cam.ac.uk

Alice Hutchings
University of Cambridge
Cambridge Cybercrime Centre
alice.hutchings@cl.cam.ac.uk

Ross Anderson
University of Cambridge
and University of Edinburgh
ross.anderson@cl.cam.ac.uk

Abstract—Legislators and policymakers worldwide are debating options for suppressing illegal, harmful and undesirable material online. Drawing on several quantitative data sources, we show that deplatforming an active community to suppress online hate and harassment, even with a substantial concerted effort involving several tech firms, can be hard. Our case study is the disruption of the largest and longest-running harassment forum KIVI FARMS in late 2022, which is probably the most extensive industry effort to date. Despite the active participation of a number of tech companies over several consecutive months, this campaign failed to shut down the forum and remove its objectionable content. While briefly raising public awareness, it led to rapid platform displacement and traffic fragmentation. Part of the activity decamped to Telegram, while traffic shifted from the primary domain to previously abandoned alternatives. The forum experienced intermittent outages for several weeks, after which the community leading the campaign lost interest, traffic was directed back to the main domain, users quickly returned, and the forum was back online and became even more connected. The forum members themselves stopped discussing the incident shortly thereafter, and the net effect was that forum activity, active users, threads, posts, and traffic were all cut by about half. The disruption largely affected casual users (of whom roughly 87% left), while half the core members remained engaged. It also drew many newcomers, who exhibited increasing levels of toxicity during the first few weeks of participation. Deplatforming a community without a court order raises philosophical issues about censorship versus free speech; ethical and legal issues about the role of industry in online content moderation; and practical issues on the efficacy of private-sector versus government action. Deplatforming a dispersed community using a series of court orders against individual service providers appears unlikely to be very effective if the censor cannot incapacitate the key maintainers, whether by arresting them, enjoining them or otherwise deterring them.

1. Introduction

Online content is now prevalent, widely accessible, and influential in shaping public discourse. Yet while online places facilitate free speech, they do the same for hate speech [1],

and the line between the two is often contested. Some cases of stalking, bullying, and doxing such as Gamergate have had real-world consequences, including violent crime and political mobilisation [2]. Content moderation has become a critical function of tech companies, but also a political tussle space, since abusive accounts may affect online communities in significantly different ways [3]. Online social platforms employ various mechanisms, for example, artificial intelligence [4], to detect, moderate, and suppress objectionable content [5], including “hard” and “soft” techniques [6]. These range from reporting users of illegal content to the police, through deplatforming users breaking terms of service [7], to moderating legal but obnoxious content [8], which may involve actions such as flagging it with warnings, downranking it in recommendation algorithms [9], or preventing it from being monetised through ads [10].

Deplatforming may mean blocking individual users, but sometimes the target is not a single bad actor, but a whole community, such as one involved in crime [11]. It can be undertaken by industry, as when Cloudflare, GoDaddy, Google and some other firms terminated service for the DAILY STORMER after the Unite the Right rally in Virginia in 2017 [12] and for 8CHAN in August 2019 [13]; or by law enforcement, as with the FBI taking down DDOS-for-hire services in 2018 [14], [15] and 2022 [16], [17], and seizing RAID FORUMS in 2022 [18]. Industry disruption has often been short-lived; both 8CHAN and DAILY STORMER re-emerged or relocated shortly after being disrupted. Police intervention is often slow and less effective, and its impact may also be temporary [11]. After the FBI terminated SIXX ROAD [19], the online drug market fragmented among multiple smaller ones [20]. The seizure of RAID FORUMS [18] led to the emergence of its successors BREACH FORUMS, EXPOSID FORUMS, and OSNI FORUMS. Furthermore, the FBI takedowns of DDOS-for-hire services cut the attack volume significantly, yet the market recovered rapidly [14], [15].

KIVI FARMS is the largest and longest-running online harassment forum [21]. It is often associated with real-life trolling and doxing campaigns against feminists, gay rights campaigners and minorities such as disabled, transgender, and autistic individuals; some have killed themselves after being harassed [22]. Despite being unpleasant and widely controversial, the forum has been online for a decade and

arXiv:2304.07037v6 [cs.CR] 6 Oct 2023



Thank you