MEMORANDUM

For Official Use Only

DATE 22 August 2018

REF [Ref]

TO [staff name and signature redacted] [redacted] 24/8/18

FROM [staff name redacted]

SUBJECT Moderation of ongoing contraventions of Trove Terms of Use

---

Background

On 16th August 2018 an interim ban was placed on the IP 101.190. [IP address partially redacted] , preventing access to Trove from that address. It was the result of multiple nuisance text corrections emanating from that address and made anonymously.

The user account *science.war.trains.etc* was the only verified activity coming from the same IP address. This user account has previously been warned about similar behaviour (see R18/76401 and R18/76874). The corrections specifically target the user account *yelnod* and are recorded in HP RM R18/76667.

The current User Moderation three-strikes Policy does not cover anonymous annotations.

The IP address ban was only moderately successful. The user has since adopted other strategies to resume their behaviour including hiding behind the anonymity of the Tor network (See *Attachment A* and *Attachment B*).

Staff have spent hours each day rolling back these nuisance corrections, managing complaints from other users and constantly updating security measures. As a result managers from the business area, IT and legal met to come up with the following mitigation strategies for both anonymous and logged in corrections that are in breach of the Terms of Use.

Immediate solutions for anonymous annotations

- **Ban the Tor Network from anonymous text correcting**
    - Prevent Trove visitors who come from the anonymity of the Tor network, from completing the captcha challenge. This prevents them from adding anonymous text corrections/tags/comments.
    - Logged in user accounts can still add annotations as normal, as can visitors who are not hiding behind the anonymity of the Tor network.

    - Visitors in the Tor network still have access to search and view collection items in Trove.
    - IT can implement this immediately and permanently.
- **Ban access to the captcha for 14 days from a known IP address when nuisance behaviour is detected**
    - As with the Tor network ban, prevent visitors from completing the captcha challenge. This time the ban is extended to a certain IP address, which has been identified as a source of nuisance behaviour.
    - This mitigation will be kept in place for 14 days.
    - Visitors from the IP address will still have access to search and view collection items in Trove, as well as contribute annotations with a user account.
    - Banning access to the captcha for a limited period ensures the next person to obtain that IP address can still visit Trove, and contribute anonymously after the ban expires. This accommodates the transient nature of IP addresses issued from residential internet providers, which are often from a shared pool, and can pass from one customer to another within days or weeks.
- **Review and roll back nuisance corrections targeted at specific users**
    - IT will produce a list of corrections that have followed targeted users.
    - Business area can review the list and will roll back nuisance corrections.
    - Review will target a defined time period –6 months prior to when the first nuisance activity was detected or reported.

Immediate solutions when correcting with a user account

- **After investigation, and verification that it is likely to be the same individual, ban all user accounts that immediately display the same behaviour**
    - Discourage users who circumvent an already issued ban and resume their activity under a different user account.
    - Take away the currency of a large corrections count associated with a user account.
    - Close the loophole that allows people to make both anonymous and identified corrections. This option also extends the existing ban to corrections made under anonymous or subsequent accounts.
- **Update the 3 strikes policy to cover anonymous user activity**
    - Bring the user moderation policy into line with the 'Blocking Trove Access' policy, which immediately bans IP addresses where no direct contact can be made with the source.
    - Note – this measure alone will not stop anonymous corrections, or Trove visitors via the Tor network

In the longer term there is a need to consider other option, including, but not limited to:

- **Authorise IT to make a decision – in consultation with the business area – on when this specific user has reappeared through other means and immediately block them.**

- Give IT and business area staff the flexibility to quickly respond, as new and unknown strategies are adopted by the user.
- Prevents further nuisance corrections which then have to be rolled back.
- Allows users such as *yelnod* to feel comfortable using Trove without being harassed.

- **Consider further mitigation strategies in the discovery system as part of TMP5: Tools for Collaborators and the Community**
  Pros:
  - Broader changes to Trove's user contribution model can be considered as part of this already scheduled Modernisation work.
  - Options could include giving users the power to choose whether their profile is public, whether Trove continues to accept anonymous annotations, whether Trove continues to have mechanisms that drive this behaviour such as the corrections leader board.
  - Would cut off all avenues for stalking and harassment.
  - Note – the timeline for such changes would be 6-9 months.

Recommendation

1. That actions 1-5 be implemented immediately, to prevent further activity from this single user.
2. That the Library undertakes a review of the Terms of Use, and the points above are considered as part of the Library's longer term strategy.

Attachment A – results of investigation into activity by IP address 101.190.[IP address partially redacted]

Examination of web logs from 1-15 August 2018 revealed this IP address is most likely a single user. Their behaviour involves visiting the Trove Forums, Newspaper zone and the profiles of a small number of users, mostly between 9pm and 2am each night. This indicates it is a single residential IP address. It minimises the likelihood they are using a shared IP address, for example at a business.

The web logs reveal this user is closely following the actions of a handful of other users (See Attachment C). Over the two week period they visited *yelnod*'s profile 133 times. They visited the user *SheffieldPark*'s profile 23 times. Both are long time Trove users.

On numerous occasions the web logs show them visiting *yelnod*'s profile, proceeding to an article *yelnod* had recently corrected, completing the captcha challenge to allow them to leave an anonymous correction, leaving a derogatory comment about *yelnod* and then moving on. They searched for the username *yelnod* as a query term 96 times over the two week period.

The requests appear to come from only two devices - A Windows 10 PC with a Chrome browser which they updated from version 67 to version 68 on 11th August at 11:30pm, and a Samsung S7 SMG930F mobile phone.

All evidence points to this being a single user. They are using Trove both anonymously and with a series of usernames created in the past three months, primarily *science.war.trains.etc* but also *corrector.namedeleted*, *Cauli.Flower*, *nlacorrectionstandards* and *radioactive.*

Attachment B – Action taken so far

Technical mitigations

The initial block implemented by NLA IT provided a 'Forbidden' http 403 warning message when the user tried to access Trove. They started using their phone connected to a mobile phone tower instead, which allowed them to circumvent the IP address block and resume their nuisance activity.

They then installed Tor on their PC – software that allows users to mask their IP address. This allowed them to circumvent the IP address block and use their PC again, continuing the nuisance activity.

Most recently it appears they have obtained a new IP address from their internet provider Telstra, which makes the initial ban ineffective and has allowed them to access Trove directly once more.

A temporary protection measure was accepted by *yelnod* – supressing their public user profile. This has mitigated the easiest method of following *yelnod*'s activity and halted the nuisance corrections for now.

However IT have observed the abuser attempting to find other avenues for following *yelnod*'s activity. Their PC has started spamming the Recent Corrections page trying to capture data.

The user is persistent and has now implemented a variety of different methods to circumvent protections the NLA IT staff can put in place. Advice from IT is that no single technical mitigation is 100% effective, including blocking a single IP address.

Staff Resources

The Trove business area has expended over 20 hours managing this single user in the past fortnight. The Branch head has spent time on top of that providing staff with advice, seeking legal advice and drafting communications. See HP RM NLA17/2393 for evidence of these correspondence and moderation activities.

NLA IT have also given significant resource, initially providing evidence of anonymous usage and then implementing the preventative measures (See Service desk Tickets #2018081510000049, #2018081610000271). They have subsequently been required to manually monitor web logs, seeking out behaviour that appears to be this user and troubleshooting to devise new protection measures as the behaviour has evolved.

Attachment C – Users followed by IP address 101.190.[IP address partially redacted] during the period 1-15 August 2018

Yelnod

1. Visited yelnod's profile 133 times
2. Searched for 'yelnod' 96 times
3. Accessed a newspaper article after searching for the term 'yelnod' 95 times

Other user profiles visited:

- SheffieldPark 23 times
- squeem 6 times
- jeri 5 times
- Gato 4 times
- Paryan 4 times
- Zebra 4 times
- Coffeepot 3 times
- MichaelP 3 times
- elmig 2 times
- radioactive 2 times
- AncestorBotherer 1 time
- Besure 1 time
- Daniela16 1 time
- Dazzling-Dropkick 1 time
- espee8800 1 time
- FrancesRaymundo 1 time
- IsambardKingdomBrunel 1 time
- Kraigbiog 1 time
- LeightonView 1 time
- llll 1 time
- MuseumofPerth 1 time

Attachment D – Browser/device information associated with IP address 101.190.[IP address partially redacted] during the period 1-15 August 2018

1. **A Windows 10 PC with a Chrome browser**
   Made 85,006 requests to Trove using the browser agent
   Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/67.0.3396.99 Safari/537.36

   The final call with this browser agent was on 11 August at 23:06.
   Then made 51,437 calls after upgrading their Chrome browser to v68 on 11 August at 23:30.
   Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/68.0.3440.106 Safari/537.36
2. **A Samsung S7 mobile SMG930F with the Samsung browser**
   516 times requests to Trove using the browser agent
   Mozilla/5.0 (Linux; Android 7.0; SAMSUNG SM-G930F Build/NRD90M) AppleWebKit/537.36 (KHTML, like Gecko) SamsungBrowser/7.2 Chrome/59.0.3071.125 Mobile Safari/537.36

Attachment E – Daily Activity by IP address 101.190.[IP address partially redacted] during the period 1-15 August 2018 10

10 August is a standout. It's the only with less than 2,000 web requests, and the only day with no access to *yelnod* related material. They loaded the Trove homepage at 5:33pm, did nothing else and left. This corresponds with no searches or access to *yelnod*'s homepage. They then came back at 11:59pm and went to the Forums, before going to recent corrections, their own user profile, and straight to the article 159061014. That article was first edited by *GJReid*, then *science.war.trains.etc* of whitespace only edits, then *yelnod*, then finally *science.war.trains.etc* at this same timestamp.

For this IP address		Web calls including a search for *yelnod*		Web calls for *yelnod*'s user profile	
Day of August	**Web calls**	**Day of August**	**Web calls**	**Day of August**	**Web calls**
1	9011	1	61	1	8
2	10214	2	49	2	7
3	4926	3	41	3	7
4	14860	4	31	4	6
5	2300	5	7	5	2
6	14304	6	62	6	8
7	2335	7	11	7	2
8	10737	8	20	8	2
9	7526	9	7	9	2
10	42	11	83	11	14
11	10184	12	56	12	14
12	9088	13	91	13	20
13	18021	14	109	14	23
14	7637	15	117	15	18
15	15774	**Grand Total**	**745**	**Grand Total**	**133**
Grand Total	**136959**				

Attachment F – Time of day Activity by IP address 101.190.[IP address partially redacted] during the period 1-15 August 2018

1. Almost half of all web calls were done between 9pm and 2am.

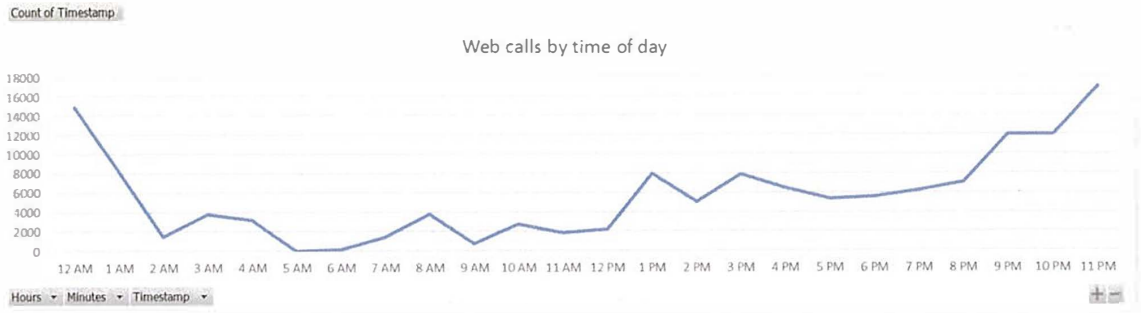

2. Web calls including *yelnod* follow a similar pattern, with activity slowing between 2am and midday.

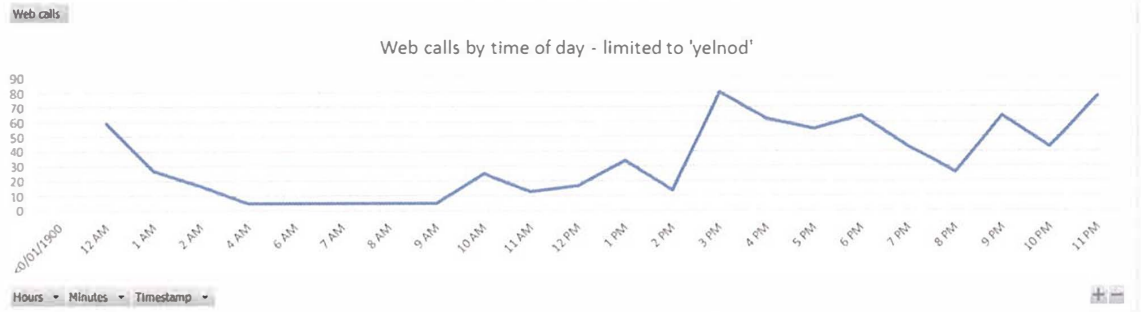

3. Web calls for *yelnod*'s user profile page have similar tendencies

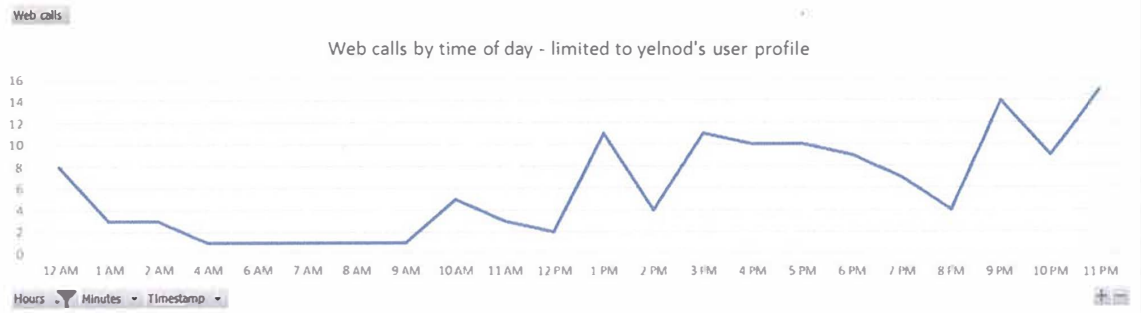

4. Web calls for completing the captcha challenge so that corrections can be made anonymously are made at similar times, particularly in the middle of the night.