

Social puncher

Twitter Purge: True Story

Part 2.

Fact check and audit
by SocialPuncher

October 2018

Twitter Purge: True Story

Part 2. Fact check and audit

In the first part of the report, we showed that the main facts of Twitter purge in July 2018 really differed from those that were covered by the media.

This coverage was mainly based on the announcement of the purge "Confidence in Follower Counts" posted on July 11 on the Investor Relations Blog blog.twtrinc.com, and on a special corporate subdomain investor.twitterinc.com, the official source for shareholders where Financial releases and Quarterly results are published.

Since the conclusion of our investigation differ significantly from the official explanation, we would like to make a fact check of key statements and a detailed audit of whole text to assess the relevance of its goal to give investors complete and reliable information about the upcoming (at that moment) purge.

We have identified 4 key statements that clearly formulated and that can be verified. But a number of important facts are missing in this text, besides the fact check, we should audit some important issues that are not included there.

Before starting the analysis, we would like to give a brief evaluation of the text of the official statement. Since there is no single form of reports on changing quantitative audience metrics for digital companies, each of them does it in its own way. First of all, such a report should be informative, it should have at least basic data.

In the Twitter announcement of the purge there is the lack of data, all dates and figures are missing completely. A number of formulations are vague and ambiguous, the text contains many generalizations and assumptions. There is not even an answer to the main question: how many accounts were locked? Many formulations are made in such a way that they can be misunderstood. In describing the main accents are often made on minor things, but the key facts are missing.

To understand what exactly the report differs from the facts revealed during the investigation, let's analyze it on specific examples. First we highlight 5 important issues not mentioned in the report, and then we check the 4 main covered facts.

5 unanswered questions about Twitter purge

What is the total number of locked accounts which were the target of this purge?

There was no Twitter statement or comment hasn't any information about it. But this is the most important issue, without its clarification, any analysis cannot be considered reliable. Obviously, there were at least 7.8 million such accounts (the number of @twitter followers, it lost after the purge). But the exact answer to this question was not given, and therefore the fact check is not possible.

Also, it was not disclosed when the very first account was locked and during which period the accounts that were affected by the purge were locked. How long these accounts have distorted the number of followers is also an important issue for understanding the significance of what happened.

How was the list of accounts created, which were removed from follower counts?

It is an unknown fact. But it certainly was not separate individual accounts that showed a sudden tweet activity. These were groups of accounts, united in botnets, managed from several control centers in different countries. These were the simplest cheap followers who throughout the whole life cycle showed no activity, except for a periodical synchronous unmotivated following. It is unclear how it was possible to detect a sudden change in their behavior, if almost $\frac{3}{4}$ of them have less than 10 tweets, and more than one third never tweeted at all. The specific metrics, exceeding the limits on which made these accounts locked, have not been disclosed.

How can you know for sure whether an account is locked or not?

3

It is absolutely impossible. In fact, this is a normal visible account that does not have any special labels. We can only determine them if you monitored accounts deleted from followers on July 12th. Moreover, after only a few days after the purge, these accounts, in spite of Twitter's statement, can be seen in the follower counts of accounts from which they were removed during the purge.

Were these accounts really locked, if we see that they are still active?

4

No one can answer this question, because, as you can see, Twitter has given us incomplete and unreliable explanations of what is actually "locked account". Therefore, there are two options for answering this question.

1. These accounts have never really been locked.
2. These accounts were locked, but their owners were able to unlock them.

How can Twitter prove that these accounts were never counted as MAU?

5

With the current disclosure of information, it cannot be proved in any way. Twitter does not publish MAU list for each quarter. Many people think that it's impossible to publish a full list of hundreds of millions of accounts, although technically it's not difficult. As a result, nobody except Twitter employees knows who is included in the MAU every quarter. The list of locked accounts is also not disclosed, as well as their exact number. It is impossible to compare the overlap of two sets without knowing what they consist of. But the most important question is how Twitter can prove that they will not be counted as MAU in Q3 2018, given that they were active during July, August, and September.

4 main Twitter's statements about the purge

1

The accounts were locked because of a sudden change in their behavior.

"If we detect sudden changes in account behavior, we may lock the account and contact the owner to confirm they still have control of it. These sudden changes in account behavior could include Tweeting a large volume of unsolicited replies or mentions, Tweeting misleading links..."

2

These locked accounts were removed from follower counts.

"This week, we'll be removing these locked accounts from follower counts across profiles globally. As a result, the number of followers displayed on many profiles may go down."

3

The accounts were owned by real people who lost control of them.

"In most cases, these accounts were created by real people but we cannot confirm that the original person who opened the account still has control and access to it."

4

The owner of the locked account has no ability to log in.

"...we've locked accounts when we detected sudden changes in account behavior [...] ...unless they validate the account and reset their passwords, we keep them locked with no ability to log in."

Investor Relations Blog

Confidence in Follower Counts

By @vijaya

Wednesday, 11 July 2018

As part of our ongoing and global effort to build trust and encourage healthy conversation on Twitter, every part of the service matters. Follower counts are a visible feature, and we want everyone to have confidence that the numbers are meaningful and accurate.

Over the years, we've locked accounts when we detected sudden changes in account behavior. In these situations, we reach out to the owners of the accounts and unless they validate the account and reset their passwords, we keep them locked with no ability to log in. This week, we'll be removing these locked accounts from follower counts across profiles globally. As a result, the number of followers displayed on many profiles may go down.

Most people will see a change of four followers or fewer; others with larger follower counts will experience a more significant drop. We understand this may be hard for some, but we believe accuracy and transparency make Twitter a more trusted service for public conversation.

Though the most significant changes are happening in the next few days, follower counts may continue to change more regularly as part of our ongoing work to proactively identify and challenge problematic accounts.

Why does an account get locked?

If we detect sudden changes in account behavior, we may lock the account and contact the owner to confirm they still have control of it. These sudden changes in account behavior could include Tweeting a large volume of unsolicited replies or mentions, Tweeting misleading links, or if a large number of accounts block the account after mentioning them. We sometimes lock an account if we see email and password combinations from other services posted online and believe that information could put the security of an account at risk — so we require accounts to change their passwords for protection. Until we confirm that everything is ok with the account, we lock it, which makes them unable to Tweet or see ads.

How are these accounts different from spam or bots?

In most cases, these accounts were created by real people but we cannot confirm that the original person who opened the account still has control and access to it.

Spam accounts (sometimes referred to as bots) typically exhibit spammy behavior from the beginning, are increasingly predictable by our systems, and can be automatically shut down with our technology. You can learn more about our ongoing work to prevent spam on Twitter [here](#).

Why just followers? Will this expand to Tweets, Likes, and Retweets?

Our ongoing work to improve the health of conversations on Twitter encompasses all aspects of our service. This specific update is focused on followers because it is one of the most visible features on our service and often associated with account credibility. Once an account is locked, it cannot Tweet, like or Retweet and it is not served ads.

Will this change affect your Monthly Active User (MAU) or Daily Active User (DAU) metrics?

No, it will not. Removing locked accounts from followers doesn't impact MAU or DAU. Locked accounts that have not reset their password in more than one month are not included in MAU or DAU. While today's change doesn't affect MAU or DAU, some accounts we remove from the service as part of our ongoing commitment to a healthy public conversation have the potential to impact publicly reported metrics.

Continuing to Communicate with You

This is another step to improve Twitter and ensure everyone can have confidence in their followers. Follow [@Twitter](#) and [@TwitterSafety](#) for the latest updates as we continue our commitment to serve the public conversation.



@vijaya

<https://blog.twtrinc.com/en/2018/Confidence-in-Follower-Counts.html>

5

Fact check of the 4 Twitter's statements about the purge

1



The accounts were locked because of a sudden change in their behavior.

FALSE

The locked accounts that were affected by the purge were mostly inactive and never changed their behavior. Up to 40% of them did not have tweets at all, 90% had less than one tweet per month on average. The only abnormal behavior that they showed was a synchronous unmotivated mass following.

2



These locked accounts were removed from follower counts.

FALSE

That was a very accurate statement before the recovery. They were only once removed from follower counts, but nothing more. They were not blocked, suspended or deleted. Now, these are regular accounts (or accounts with hidden following numbers), that counted as followers again.

3



The accounts were created by real people who lost control of them.

FALSE

No, these accounts were never owned by real people. These are primitive bots, created only for sale as cheap followers. Their only goal is to simulate Twitter statistics. Most of them follow hundreds of completely different accounts, which have mostly purchased followers.

4



The owner of the locked account has no ability to log in.

FALSE

A few days after the purge, millions of "locked accounts" synchronously performed the same action: first, they unfollowed all, and then follow some former friends again. As a result, hundreds of accounts have returned their audience, at least ten accounts have restored more than a million followers. The owners of these botnets still have full access to the locked accounts.

Conclusion on the fact check and audit



"... we believe accuracy and transparency make Twitter a more trusted service for public conversation."
"Confidence in Follower Counts"

About accuracy

After the fact check, it was found out that 4 of the key 4 statements were false, and three of them could be refuted before purge, after analyzing the list of accounts to be removed from follower counts.

About transparency

5 important questions about the purge were't answered. The report tries to avoid them as much as possible, some of them completely ignored, the information about others is incomplete. Those issues that could not be avoided are presented in the report as an unsubstantiated statement, without any evidence that can be verified. Here are 4 key metrics that should be disclosed first for the real transparency.

1. The exact number of so-called "locked accounts", removed from follower counts.
2. Their basic stats: the average number of tweets, followers, and followees.
3. Monthly stats on the number of locked accounts since the first one was locked.
4. All specific reasons why these accounts were locked and their proportion.

Only the disclosure of data on these 4 items can confirm the accuracy of the MAU numbers in the 2018 Q3 report.