

Users of Snowflake as a pluggable transport were invited to submit their feedback from 24 February – 30 March 2021.

Summary

The Tor Project surveyed users from across the world about their experience using Snowflake as a client. Of the 1,796 participants who submitted the survey, 726 confirmed they use Snowflake as a pluggable transport. The typical Snowflake user has less than a year’s experience with Tor Browser, connects to Snowflake from their mobile phone, and has tested Snowflake for ~1 hour in a single session.

Although many Snowflake users experienced some degree of slowness while connecting and/or browsing, only between one-quarter and one-third of participants found this completely prohibitive. Similarly, around three-quarters of users would recommend Snowflake in its current form.

Background

Starting in March 2021, the Tor Project began a push via community channels to increase the number of Snowflake proxies and users of Snowflake as a pluggable transport ahead of its planned stable release in Tor Browser 10.5. This push served three main purposes:

- To stress test the system, and monitor potential resulting slow-downs, server issues or new bugs.
- To identify any edge case bugs that haven’t been reported or addressed yet.
- To learn more about how the bandwidth and latency of Snowflake differs depending on geographical location.

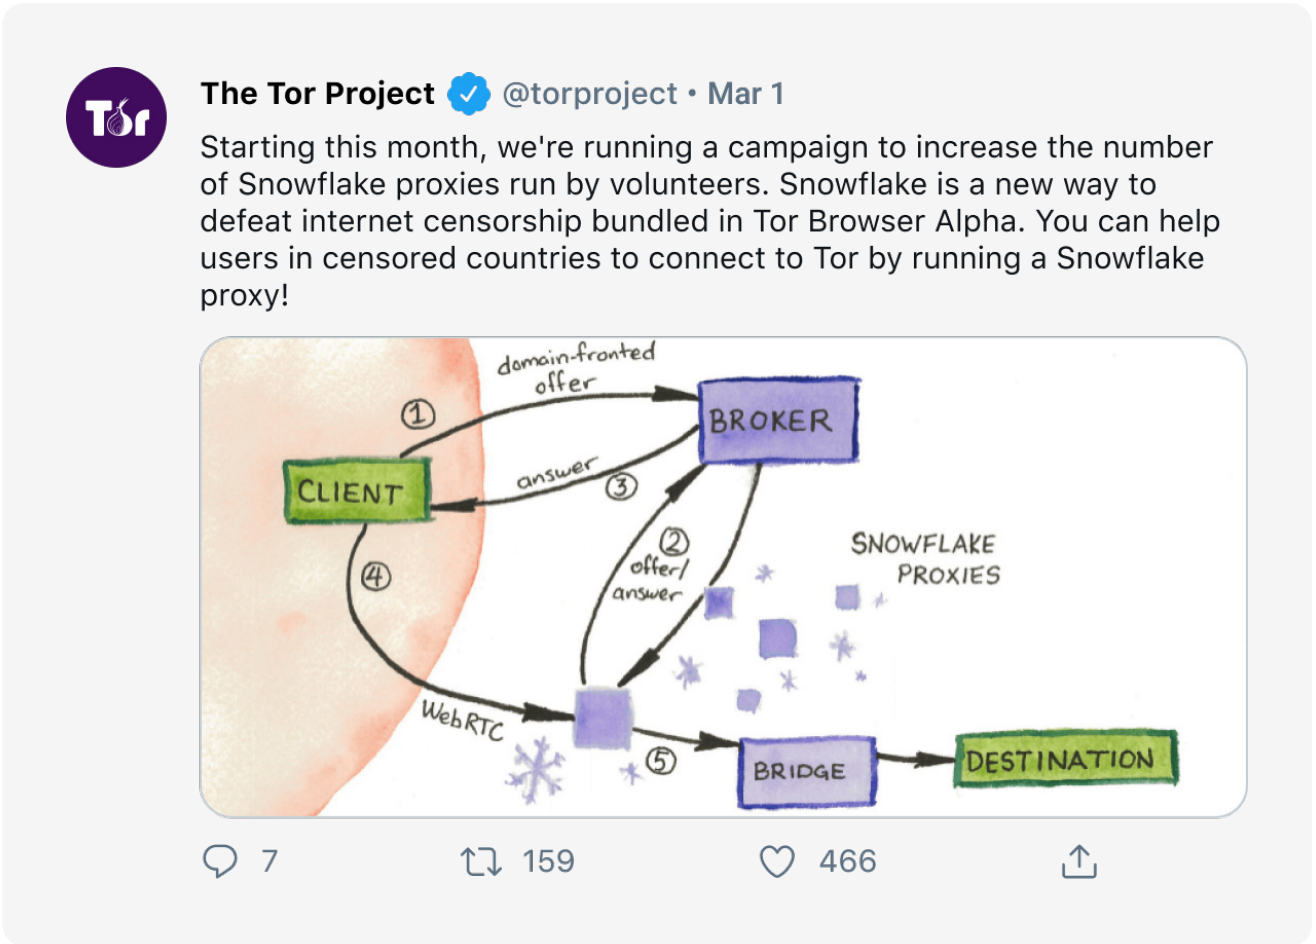


Fig 1: Snowflake promotional tweet

Methodology

In addition to the platform measurements typically available to the Tor Project, the Anti-Censorship Team required an avenue for users to report back through during the stress test – and provide more qualitative feedback about their browsing experience and connection quality while using Snowflake.

As such, the team designed and built an anonymous 14 question survey, hosted on the Tor Project’s self-hosted LimeSurvey instance, available on both survey.torproject.org and as an Onion Service. Running from 24 February – 30 March 2021, participants were recruited through two principal avenues:

- Via the Tor Project’s social media channels, primarily Twitter, and mailing lists.
- Via a banner deployed on Tor Browser Alpha’s homepage (about:tor) for the duration of the survey.

Volunteers who run Snowflake as a proxy were specifically discouraged from participating. Rather, recruitment focused on soliciting the feedback of users of Snowflake as a client.

Lastly, results were analyzed using a locally-hosted instance of the open source database visualization tool Metabase.

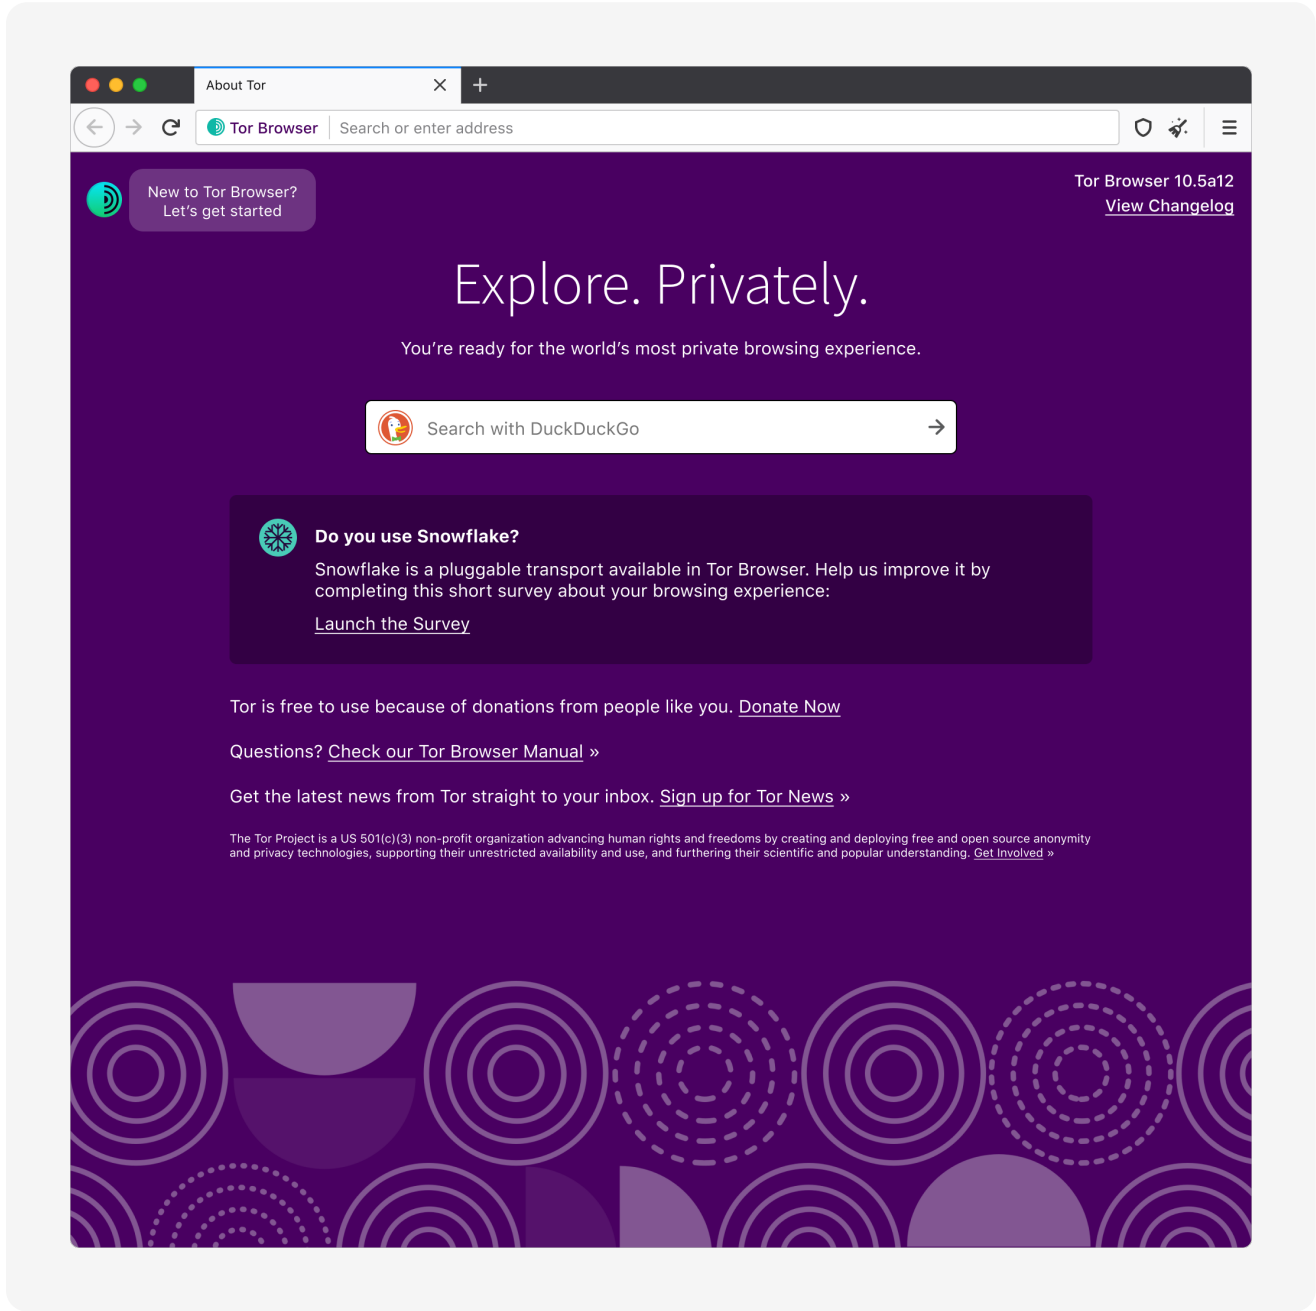


Fig 2: Snowflake banner on about:tor



The survey received 1,795 complete responses, of which 726 participants confirmed they use Snowflake as a pluggable transport.

Sample size and limitations

The survey received 1,797 submissions, out of 8,660 total responses (counting partial attempts and abandonments). Due to the anonymous nature of the survey, the total number of responses includes a degree of spam and duplicate, though not identical, submissions (i.e. participants who submitted the survey more than once).

The banner on Tor Browser Alpha was the top referrer to the survey, as self-reported by participants. Due to the broad audience reached by our recruitment methods, participants were required to confirm their use of Snowflake in order to filter erroneous responses. Thus, all subsequent statistics and visualizations feature confirmed Snowflake users exclusively.




Fig 3: Top referrers (note: “Other” includes the Tor Blog)


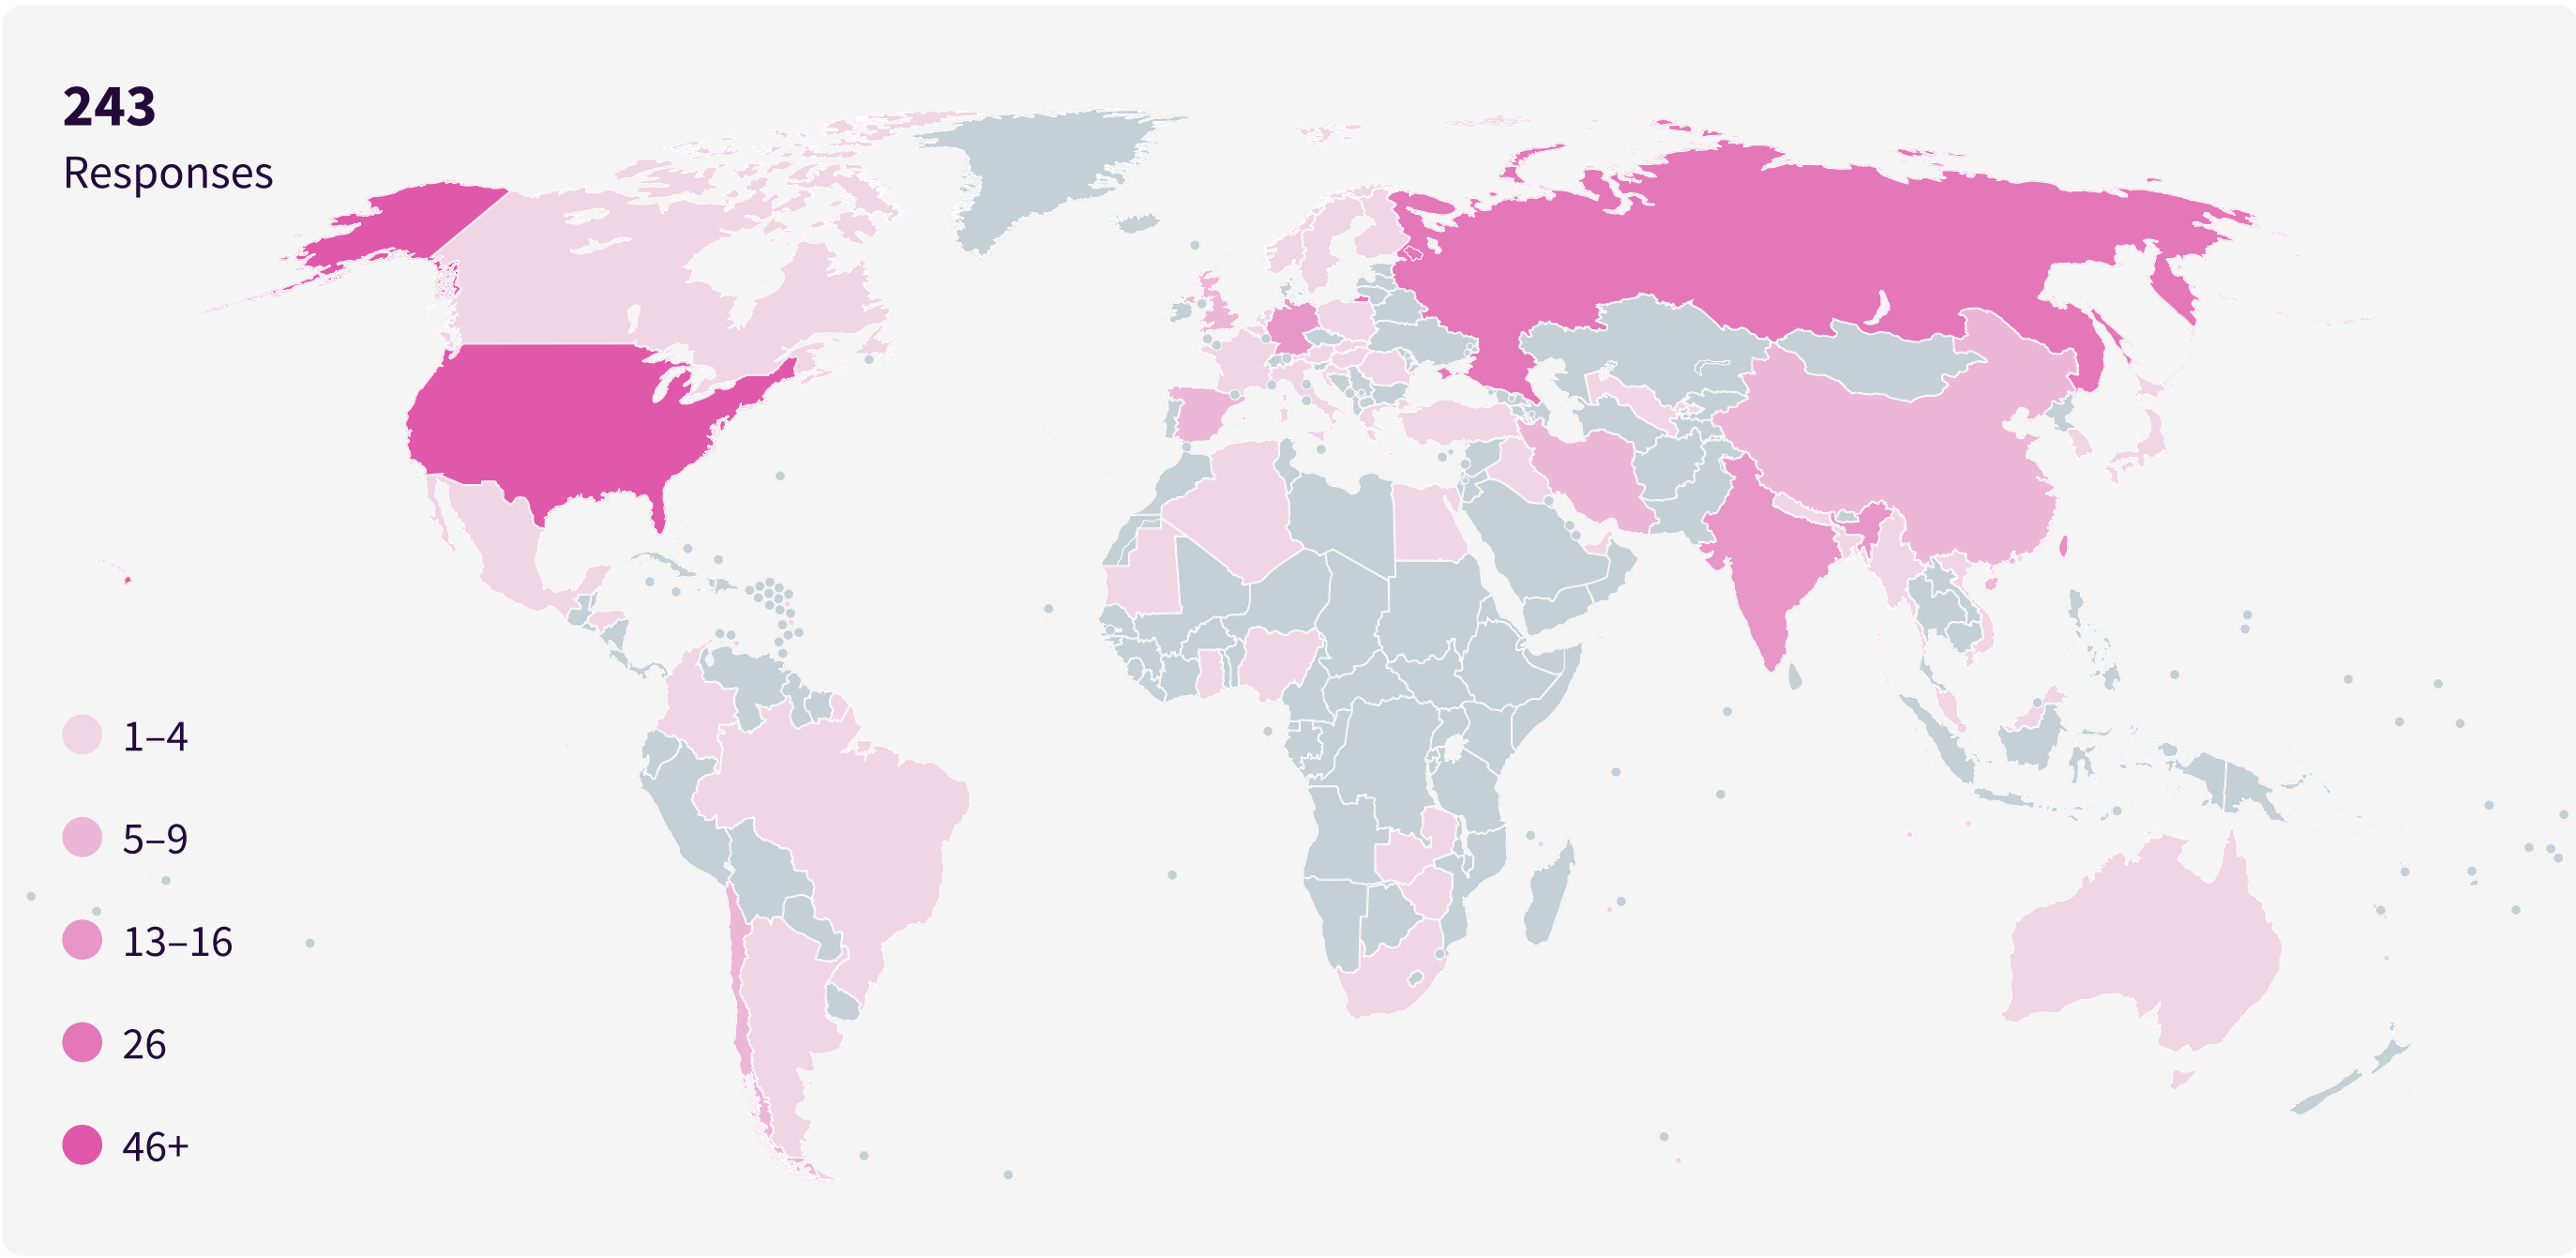


Fig 4: Do you use Snowflake as a Pluggable Transport?



North America +5 South America +1 Caribbean +1 Europe +1 Middle East +1 Africa +1 CIS +1 Asia +1

Fig 5: Snowflake users by given country, including the number of additional participants who responded by region instead



The majority of Snowflake users who completed the survey began using Tor Browser several times a week within the past year.

Tor Browser usage patterns

243 participants (constituting 33% of confirmed Snowflake users) volunteered the location they were connecting from, which was marked as optional in order to protect the privacy of high-risk users. Of those, 231 provided their location at the country-level, whereas 12 opted to provide a general region instead. The top five locations with the highest number of Snowflake users includes:

- United States 46
- Russia 26
- India 16
- Germany 13
- Mainland China 9

However, anglophone and relatively censorship-free locations may be overrepresented due to several factors: the survey and recruitment were conducted in English-only; participants from high-risk countries may be significantly less likely to volunteer their location; and accessing the survey will naturally be more difficult, if not impossible, in locales where Tor, Snowflake and/or the torproject.org website are blocked.

Participants tended to be relatively new to Tor Browser, with those having used Tor Browser for less than a year constituting 55% of confirmed Snowflake users. Additionally, the majority of Snowflake users use Tor Browser near daily – with notably frequent use observed among new (i.e. <6 months) and longtime (i.e. ≥6 years) users in particular.




Fig 6: Frequency grouped by total time using Tor Browser (note: due to a regrettable error, 3-5 years should read 3-6 years)



Despite frequent use of Tor Browser, the majority of confirmed Snowflake users have only used Snowflake as a client lightly.

Snowflake usage patterns

70% of confirmed Snowflake users connect to Snowflake via Tor Browser on mobile devices (e.g. Tor Browser for Android), with only 19% connecting from Tor Browser on desktop.

Participants’ self-reported time spent using Snowflake, total number of Snowflake sessions and longest snowflake session each result in extremely positively skewed distributions when plotted – to the extent that any visualization borders on illegible. This is due to the modal value for each being 1, despite the extreme ranges observed: 599.9, 299 and 199.9 respectively.

Approximately 63% of confirmed Snowflake users have used Snowflake for 5 hours or fewer. ~61% have used Snowflake for 5 or fewer sessions, and ~71% of users’ longest session did not exceed 5 hours.




Fig 7: Snowflake users by platform

Snowflake browsing experience

69% of confirmed Snowflake users reported some degree of increased waiting time during bootstrapping, however only 30% indicated that the longer bootstrapping times would discourage them from using Snowflake again. Similarly, 65% of confirmed users reported slower speeds when connecting via Snowflake than they would normally experience, however only 25% found the speed completely prohibitive.

The same ratios roughly correlate to the number of confirmed Snowflake users who reported issues with Snowflake after configuration, at 35%, and the number of users who are either unsure or would not recommend Snowflake to their peers, at 25%. However, figures 9 and 10 are approximate at best due to the methods used to analyze and quantify participants’ responses, which were submitted in free-text form.




Fig 9: Proportion of users who experienced issues

		Lowest	Highest	Mode	Average	Total
Time spent using Snowflake	471	5 mins	500 hours	1 hour	9.83 hours	4,630 hours
Number of Snowflake sessions	459	1 session	300 sessions	1 session	6.44 sessions	2,857 sessions
Longest Snowflake session	434	5 mins	200 hours	1 hours	5.94 hours	2,579 hours

Fig 8: Time spent using Snowflake, number of Snowflake sessions and longest reported Snowflake session



75% of users had a positive view of Snowflake, although many experienced connection troubles and slow speeds while browsing.

Of the 726 responses submitted for each, approximately 230 to 260 were found to be uncountable using simple string-matching methods. Additionally, many participants entered random characters in order to skip the question(s) – rendering their responses incomprehensible. A full list of genuine responses over 15 characters in length has been included for reference in the non-public version of this report.

To identify the locations in which Snowflake is reported to perform poorly, the proportion of confirmed Snowflake users who selected either “Too long/slow” or “Much longer/slower” in response to questions about Snowflake’s bootstrapping duration and browsing speeds are illustrated in figure 11. To achieve some level of confidence in the results, only countries with ≥5 confirmed Snowflake users have been counted, reducing the number of countries featured from 55 to 12.




Fig 10: Proportion of users who recommend Snowflake

Country	Freq.	Long Bootstrap	Slow speed
Chile	5	60%	20%
Russia	26	50%	46.2%
Belarus	7	57.1%	57.1%
Germany	13	53.8%	30.8%
India	16	50%	12.5%
Iran	8	50%	37.5%
Russia	26	50%	46.2%
China	9	33.3%	22.2%
Spain	6	33.3%	33.3%
US	46	21.7%	26.1%
UAE	6	16.7%	16.7%
UK	6	0%	0%

Fig 11: % of participants by country (with ≥5 responses) who experienced significant slowness




Fig 12: Perception of Snowflake’s bootstrapping vs browsing speeds, relative to not using Snowflake