



ROBOTROLLING

PREPARED AND PUBLISHED BY THE
**NATO STRATEGIC COMMUNICATIONS
CENTRE OF EXCELLENCE**



Executive Summary

This quarter, we observed a significant drop in both authentic and inauthentic engagement with the topic of NATO in Poland and the Baltics. The number of bots and the volume of messages they disseminated decreased among both English- and Russian-language communities on Twitter and VK.

Spikes in bot activity this winter coincided with NATO military exercises. Inauthentic accounts placed particular emphasis on unfounded claims of disorderly conduct among NATO soldiers and the alleged effects that military exercises have on local civilian populations. Throughout this period, inauthentic accounts also amplified claims of turmoil within the alliance and fears of military buildup along the border of Kaliningrad.

While Russian-language bot activity was focused primarily on military affairs, English-language bot activity was centered on US

affairs in the wake of the 2020 presidential election, particularly debates over how the Biden administration will impact US-Polish relations and transatlantic security more generally.

In this issue of *Robotrolling*, we also discuss the steps Twitter has taken to protect its platforms from attempts to incite violence, organise attacks, and share misinformation following the riot at the US Capitol on 6 January 2021. This regulatory enforcement resulted in the removal of tens of thousands of accounts connected to QAnon conspiracy theorists. Our analysis is accompanied by a visualisation of the English-language accounts mentioning the NATO presence in Poland and the Baltics, demonstrating the impact these account removals will have on the information space. ■

The Big Picture

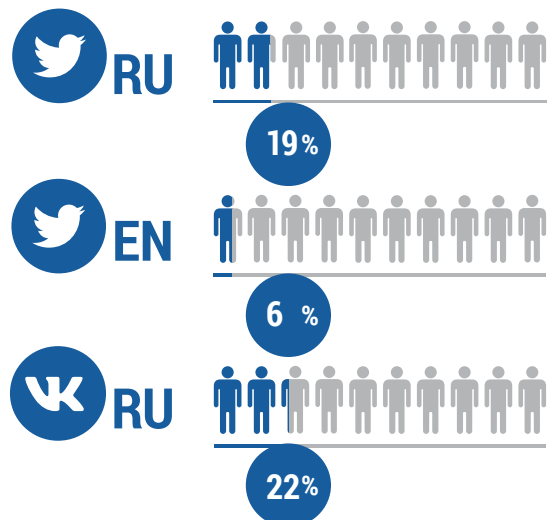
This edition of *Robotrolling* continues to monitor the online manipulation of information regarding the NATO presence in Poland, Estonia, Latvia, and Lithuania on the social media platforms Twitter and VK. Our analysis focuses on the activities of automated accounts (bots) and coordinated, anonymous human accounts (trolls). This issue tracks the key trends that emerged in the Russian- and English-language information spaces during the period 1 November 2020 to 31 January 2021.

This quarter, the English- and Russian-language information spaces were discernibly quiet. We noted that the overall volume of messages referencing the Enhanced Forward Presence (eFP) dropped significantly compared to the previous period, 1 August to 31 October 2020. On Twitter, we observed a total of 5 300 messages about the eFP, which is roughly half the output recorded last quarter. The number of unique users engaging in discussions about NATO also shrank by a similar margin. This drop is the result of a decrease in both authentic and inauthentic activity.

In sum, bot messaging on Twitter decreased by nearly one third throughout the end of 2020 and first month of 2021. The volume of English-language bot activity dipped by 20%, while the number of active bot users was less than half that of the previous quarter. The share of conversations driven by English-language bots dropped to a mere 8%.

Last quarter, we observed a reduction in the volume of Russian-language bot messages and the number of accounts disseminating them on Twitter. This quarter, the absolute numbers were even further reduced by more than a third. However, due to accompanying decreases in recognisably human and anonymous Russian-language activity, the portion of activity attributed to bots rose from 22% to 27% of messages. In other words, bot activity commanded slightly more of the information space because rates of anonymous and human activity fell.

Automated activity on VK this quarter mirrored our findings on Twitter. The total volume of VK posts decreased from nearly 12 000 to 7 600—a 35% drop—while bot posts decreased by a third, from 3 700 to 2 600 messages. The number of both authentic and inauthentic users fell by a quarter. In a similar pattern to that observed on Twitter, the portion of activity attributed to bots on VK increased marginally as other types of engagement waned. ■



Country Overview

Although bot activity this quarter was low, we observed several peaks in messaging driven by inauthentic accounts on both Twitter and VK. On Twitter, Russian-language bot activity peaked on 2 November, at the very beginning of our monitoring period, coinciding with the annual Steel Shield exercise in the Baltics. The most significant spike on VK was driven by a single user who reposted the same incoherent message, which is irrelevant to the eFP, over one hundred times. In the English-language Twitter space, bot activity culminated on 9 January as inauthentic accounts amplified former US Secretary of State Mike Pompeo's announcement of the signing of an Enhanced Defence Cooperation Agreement with Poland.

Throughout this quarter, escalations in Russian-language bot activity corresponded with NATO military exercises, fears of military buildup near Kaliningrad, and rumours of turmoil within the alliance. In contrast, inauthentic English-language conversations revolved around the 2020 US presidential election and debated how the incoming Biden administration would impact transatlantic security.

Estonia

This quarter, Russian-language bots latched onto a specific element in a controversy over the marriage referendum bill proposed by a populist party. In order to slow the progress of the bill, which would call for a referendum on whether marriage should be defined as being between a man and a woman, opposing parties added thousands of amendment proposals as an obstructionist tactic. Amongst these amendments, proposals to hold referendums on Estonia withdrawing from NATO and restoring the Soviet Union were gleefully promoted by Russian-language media. Inauthentic accounts circulated these stories, which portrayed the questions as emblematic of weakness in both Estonian leadership and the NATO alliance.

Latvia

Latvia received the second highest level of Russian-language bot attention but simultaneously the fewest mentions by English-language bots. Following the Riga Conference in November, Russian-language bots seized upon statements made by Latvian PM Kariņš and NATO Secretary General Stoltenberg describing Russia as posing a security threat to its regional neighbours. Throughout December, bot activity rose with reports of Latvian projects to modernise infrastructure that would better accommodate NATO forces. In late January, inauthentic accounts mocked Latvia for donating British off-road vehicles allegedly developed in the 1940s to the Ukrainian Armed Forces.

Lithuania

Lithuania received the most attention in inauthentic Russian-language content this quarter. Lithuanian affairs were targeted by several articles from the pro-Kremlin outlet *Baltnews*, which were further amplified by bot accounts. In early November, automated accounts shared a *Baltnews* article that referred to Iron Wolf 2020 exercises as a "hotbed" of COVID-19 infection. Activity targeting Lithuania spiked again in late January, when bots circulated a *Baltnews* article accusing the government of prioritising military spending during the COVID-19 pandemic.

Poland

Poland remained the country most discussed by English-language bots and returned to being the least-referenced country by Russian-language bots. English-language bot messaging was not driven by any singular event, except for outgoing Secretary of State Mike Pompeo's reflections on the signing of the Enhanced Defence Cooperation Agreement with Poland. Instead, inauthentic accounts discussed the outcome of the US election. Russian-language bots amplified one fake story, claiming that unidentified 'Polish experts' are wary of transferring the command of joint operations to Turkey. These reports claimed that there is distrust of Turkey within NATO that is causing cracks to form between allies. ■

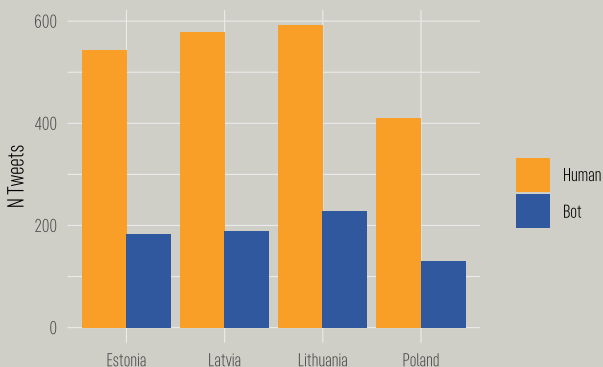


Figure 2: Distribution of Russian-language tweets mentioning NATO and Estonia, Latvia, Lithuania, or Poland




Figure 3: Percentage of bot-generated Twitter messages in English and Russian over the past three quarters

Themes

In the previous edition of *Robotrolling*, we observed that discussions about the eFP were largely dominated by the 2020 presidential election in Belarus. In this edition, although we found that Belarusian narratives continued to be circulated by some anonymous users, we noted that bot users appeared to have moved on to a new topic. This quarter, spikes in bot activity repeatedly coincided with NATO military exercises and alleged NATO tensions with Russia and within the alliance itself. In coverage of and discussions about NATO exercises, we observed the resurgence of a common narrative used to discredit the eFP: NATO troops stationed in the Baltics and Poland are dangerous, destructive, and unwelcome occupiers. Additionally, these host countries are portrayed as ignoring their citizens' discomfort with the NATO presence, essentially prioritising the needs of the alliance over those of their respective populations.

The theme of disorderly soldiers first emerged on 2 November, the peak of Russian-language bot activity, in reports of the NATO Steel Shield exercises being held in the Baltics. Various Russian-language media outlets shared the same story that local residents complained about the behaviour of military personnel, claiming they insult the local population, urinate on monuments, and engage in drunken brawls. In addition to their behaviour, NATO soldiers allegedly cause damage to cities and accidents with NATO equipment. The only evidence of these incidents provided in the articles takes the form of quotes attributed to unspecified local residents.

This narrative thread was woven into coverage of other NATO military exercises. Bot discussions portrayed the Iron Wolf exercises in Lithuania, also held in early November, as impeding civilian life. Reports shared unfounded claims that the population does not approve of military exercises on Lithuanian territory, particularly because the movement of NATO equipment causes traffic jams. In their coverage of the Tumas 2020 exercises, Russian- and Polish-language media circulated claims that American soldiers had engaged in a foul-mouthed brawl with their Polish allies, and subsequently blamed Polish authorities for the incident; membership in NATO was framed as an 'overlord-vassal' relationship.

Descriptions of plans to modernise infrastructure in the region furthered this characterisation of Poland and the Baltics as submitting to NATO occupation. In early December, reports of a project to renovate Riga airport, which would include measures to accommodate NATO air forces in the future, drew criticism as wasting national funds and 'serving' NATO. Later that month, news that the Latvian Defense Ministry plans to reconstruct a training ground near Daugavpils was accompanied with accusations that government authorities are ignoring local concerns over the project. Overall, bots this quarter disseminated a distorted image of NATO as terrorising the member-state populations and monopolising the financial budgets of its allies. ■




Figure 4: Timeline showing spikes in activity about the NATO presence on English and Russian-language Twitter, as well as VK.

Twitter's QAnon Crackdown

In January 2021, Twitter announced it had removed more than 70 000 accounts associated with the QAnon conspiracy theory movement. The crackdown followed the riot of 6 January 2021, when a belligerent crowd stormed the US Capitol building in an attempt to disrupt the outcome of the 2020 US presidential election. QAnon believers featured among the crowd that forced its way into the heart of the US Congress. According to Twitter, the deleted accounts were 'engaged in sharing harmful QAnon-associated content at scale and were primarily dedicated to the propagation of this conspiracy theory across the service' that increased the risk for offline harm as well.

The QAnon movement originated in 2017 as a conspiratorial pro-Trump movement against the so-called 'deep state' but has been consistently creeping into mainstream conservative politics. QAnon is an internet conspiracy whose adherents believe that the world is secretly run by Satan-worshipping Democrats, pedophiles, and human traffickers. QAnon supporters artificially inflate hashtags and coordinate online abuse of perceived foes, including politicians.

To better understand the effects of Twitter's recent enforcement action, we mapped out the English-language ecosystem monitored by Robotrolling. Figure 5 is a network graph showing connections between English-language accounts that mentioned NATO in the Baltics and Poland in the previous 12 months. Accounts deleted in early January 2020 are marked

in black. The large green cluster consists of conservative voices, QAnon-supporters, and pro-Trump accounts. The large orange-cluster contains liberal voices and anti-Trump accounts. The large blue area comprises accounts associated with promoting democracy, especially in Hong Kong and Belarus. The large pink area in the centre of the graph is made up of accounts without a singular focus; however, their messaging tends to concentrate on international relations, news, and the military. The smaller reddish-pink clusters consist of users who primarily discuss national politics in Poland, the UK, and Ukraine.

Our analysis confirms that the vast majority of the recently deleted accounts are associated with QAnon. Prior to their removal by Twitter, QAnon supporters also engaged in discussions about the NATO presence in the Baltics and Poland. In the past we have seen these accounts driving four principal spikes, all of which have been discussed in previous Robotrolling reports:

- 1) promoting tweets by the US embassy in Poland on relocating US troops away from Germany
- 2) retweeting comments by Brad Parscale about NATO
- 3) trolling Representative Adam Schiff in February 2020
- 4) amplifying reports of Biden's comments about Hungary and Poland in October 2020



Figure 5: Deleted accounts (black) in English-language Twitter

Twitter's QAnon Crackdown (continued)

Apart from these spikes in coordinated messaging, QAnon accounts have featured sporadically in our data. When they were active, these accounts primarily retweeted messages, sometimes at very large volumes. The QAnon cluster featured a mixture of automated accounts, accounts held by recognisable humans, and anonymous accounts. The largest group of deleted accounts were classified as anonymous accounts by our algorithm, indicating that they were operated by people. According to Twitter's comments about the removals, many individuals impacted by permanent suspension had been operating multiple QAnon-affiliated accounts.

On the whole, we found that these accounts were not generally talking about NATO, instead focusing their efforts on promoting President Trump and attacking his enemies. Thus when an elite voice within this community happened to mention NATO, we observed a large spike in activity—not because they were particularly focused on the alliance, but because they systematically amplified all messaging. This is reflected in the network map, which shows comparatively few deleted accounts in the central pink area where the majority of conversation about the Alliance occurs. While the removal of potentially harmful QAnon accounts will provide some protection from misleading information in Twitter conversations moving forward, their removal will likely not have a significant impact on inauthentic networks discussing NATO's presence in Poland and the Baltics. ■

Prepared by Dr. Rolf Fredheim and Kristina Van Sant, published in March 2021 by

**NATO STRATEGIC COMMUNICATIONS
CENTRE OF EXCELLENCE**

The NATO StratCom Centre of Excellence, based in Latvia, is a multinational, cross-sector organization which provides comprehensive analyses, advice and practical support to the Alliance and Allied Nations.

www.stratcomcoe.org | [@stratcomcoe](https://twitter.com/stratcomcoe) | info@stratcomcoe.org