


Google Australia Pty Ltd



www.google.com.au

Google Australia thanks the Select Committee on Foreign Interference through Social Media for the opportunity to respond to this Inquiry [on the risk posed by foreign interference through social media](#).

The Internet has enabled people to create, connect, and distribute information in new and innovative ways. It has exposed us to perspectives and experiences that were previously out-of-reach. It has enabled increased access to knowledge for more people than ever before.

Google continues to believe that the Internet offers significant benefits to society – contributing to global education, healthcare, research, and economic development by enabling citizens to become more knowledgeable and involved through access to information at an unprecedented scale. However, like other communication channels, the open Internet is vulnerable to the organised propagation of false or misleading information.

These concerns directly affect Google and our mission – to organise the world’s information and make it universally accessible and useful. When our services are used to propagate deceptive or misleading information, our mission is undermined.

How companies like Google address these concerns has an impact on society and on the trust users place in our services. We take this responsibility very seriously. We believe that meeting it begins with providing transparency into our policies, inviting feedback, enabling users to understand and control their online engagement, and collaborating with policymakers, civil society, and academics around the world in the development of sensible, effective policies and processes.

This submission intends to provide a brief overview of the work Google has done to counter coordinated influence operations and other government-backed attacks. For further detail on how Google prevents disinformation across all its platforms, we recommend reference to [How Google Fights Disinformation](#). We hope this work supports the Select Committee to achieve its goals.

## Google’s approach to tackling disinformation

We have an important responsibility to our users and to the societies in which we operate to curb the efforts of those who aim to propagate false information on our platforms. At the same time, we respect our users’ fundamental human rights (such as free expression) and we try to be clear and predictable in our efforts, letting governments, users and content creators decide for themselves whether we are operating fairly. Of course, this is a delicate balance as sharing too many granular details of how our systems and processes work would make it easier for bad actors to exploit them.

Algorithms cannot determine whether a piece of content on current events is true or false, nor can they assess the intent of its creator just by reading what's on a page. However, there are clear cases of intent to manipulate or deceive users. For instance, a news website that alleges it contains "Reporting from Canberra, Australia" but whose account activity indicates that it is operated out of Eastern Europe is likely not being transparent with users about its operations or what they can trust it to know firsthand. To address this situation, our policies across Google Search, Google News, YouTube, and our advertising products clearly outline behaviors that are prohibited – such as misrepresentation of one's ownership or primary purpose on Google News and our advertising products, or impersonation of other channels or individuals on YouTube.

Government-backed or State-sponsored groups who attempt to gain access to our user's accounts have varying goals in carrying out operations targeting Google's products: Some are looking to collect intelligence or steal intellectual property; others are targeting dissidents or activists, or attempting to engage in coordinated influence operations and disinformation campaigns. Our products are designed with robust built-in security features, like Gmail protections against phishing and Safe Browsing in Chrome, but we still dedicate significant resources to developing new tools and technology to help identify, track and stop this kind of activity as it evolves. In addition to our internal investigations, we work with law enforcement, industry partners, and third parties like specialised security firms to assess and share intelligence.

## Coordinated Influence Operations

Our work tackling disinformation is an important pillar of a wider, holistic and years-long effort to tackle information threats: for many years now, we have invested heavily to counter efforts seeking to deceive, harm, or take advantage of users, including by developing industry-leading technology to protect our users against spam, malware, and "content farms". This includes countering targeted and government-backed operations against Google and our users.

We continue to communicate our findings on government-backed phishing, threats and disinformation, and our Threat Analysis Group has recently launched a [new quarterly bulletin](#) to share information about actions we take against accounts that we attribute to coordinated influence campaigns. For the first quarter of 2020, we reported disabling influence campaigns originating from groups in Iran, Egypt, India, Serbia and Indonesia<sup>1</sup>. Since March, we've removed more than a thousand YouTube channels that were apparently part of a large campaign and that were behaving in a coordinated manner. These channels were mostly uploading spammy, non-political content, but a small subset posted primarily Chinese-language political content supporting Chinese Communist Party (CCP) policy and propaganda positions, similar to the findings of a recent Graphika report<sup>2</sup>.

On any given day, Google's Threat Analysis Group (TAG) is tracking more than 270 targeted or government-backed attacker groups from more than 50 countries<sup>3</sup>. These groups have

---

<sup>1</sup> Google Threat Analysis Group Blog - [TAG Bulletin: Q1 2020](#)

<sup>2</sup> Graphika report - [Return of the \(Spamouflage\) Dragon](#)

<sup>3</sup> Google Threat Analysis Group Blog - [Updates about government-backed hacking and disinformation](#)


different goals in carrying out their operations: while security attacks may focus on collecting intelligence or stealing intellectual property, coordinated influence operations and disinformation campaigns may be financially motivated, engaging in disinformation activities for the purpose of turning a profit; others are politically motivated, engaging in disinformation to foster specific viewpoints among a population, to exert influence over political processes, or for the sole purpose of polarising and fracturing societies.

When we detect attempts to conduct coordinated influence operations on our platforms, whether state-backed or otherwise we swiftly remove such content from our platforms and terminate these actors' accounts. We take steps to prevent possible future attempts by the same actors, and routinely exchange information and share our findings with others in the industry.

## Lessons from Targeted Campaigns

Manipulation of information services for political or ideological aims is not limited to networks of inauthentic accounts on social channels attempting to influence genuine users. Targeted attacks on individuals and organisations have continued, with their methodology changing as Google and other digital platforms refine protective strategies.

Which groups are targeted by government-backed attackers will be of no surprise to the committee. Government-backed attackers consistently target geopolitical rivals, government officials, journalists, dissidents and activists. For example, the chart below details the Russian threat actor group SANDWORM's targeting efforts (by sector) over the last three years<sup>4</sup>.



Distribution of targets by sector by the Russian threat actor known as SANDWORM

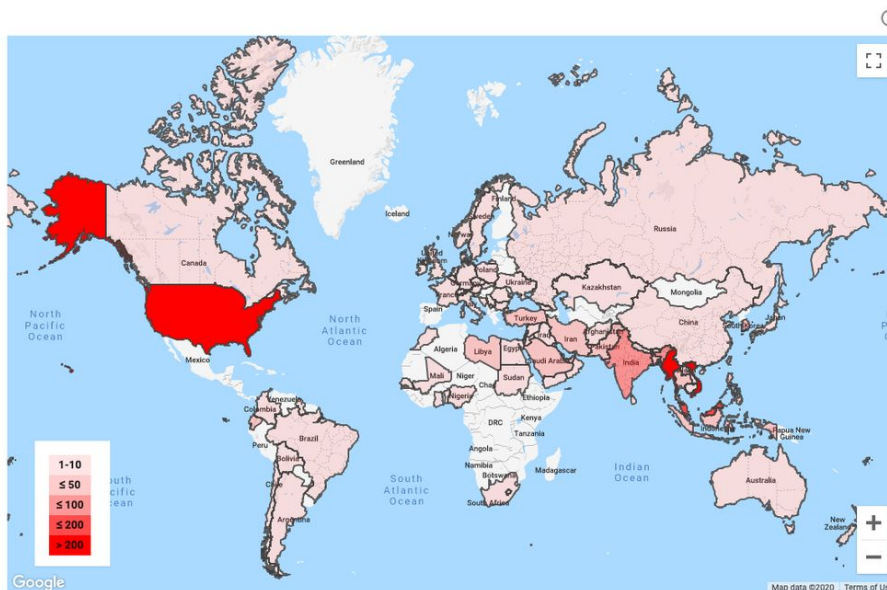
<sup>4</sup> Google Threat Analysis Group Blog - [Identifying vulnerabilities and protecting you from phishing](#)

Government-backed attackers also tend to repeatedly attack their targets. In 2019, one in five accounts that received a warning was targeted multiple times by attackers. If at first the attacker does not succeed, they'll try again using a different lure, different account, or trying to compromise an associate of their target.

Since the beginning of 2020, we've seen a rising number of attackers, including those from Iran and North Korea, impersonating news outlets or journalists. For example, attackers impersonate a journalist to seed false stories with other reporters to spread disinformation. In other cases, attackers will send several benign emails to build a rapport with a journalist or foreign policy expert before sending a malicious attachment in a follow up email. Government-backed attackers regularly target foreign policy experts for their research, access to the organisations they work with, and connection to fellow researchers or policymakers for subsequent attacks.

In April 2020 alone we sent 1,755 warnings to users whose accounts were targets of government-backed attackers<sup>5</sup>.

We intentionally send warnings in timed batches to all users who may be at risk, rather than at the moment we detect the threat itself, so that attackers cannot track some of our defense strategies. We also notify law enforcement about what we're seeing, as they have additional tools to investigate these attacks.



Distribution of the targets of government-backed phishing attempts in April 2020

Separately, we provide Google's [Advanced Protection Program](#) (APP) to journalists, government officials, human rights advocates and others who may be at high risk. We have yet to see people successfully phished if they participate in Google's APP, even if they are repeatedly targeted. Our APP provides the strongest protections available against phishing and account hijacking and is specifically designed for the highest-risk accounts.

<sup>5</sup> Google Threat Analysis Group Blog - [Updates about government-backed hacking and disinformation](#)

## Case Study: COVID-19

Bad actors frequently look at crises as an opportunity, and COVID-19 provides one such example. Across Google products, we're seeing bad actors use COVID-related themes to create urgency so that people respond to phishing attacks and scams. Our security systems have detected examples ranging from fake solicitations for charities and NGOs, to messages that try to mimic employer communications to employees working from home, to websites posing as official government pages and public health agencies.

Recently, our systems detected 18 million malware and phishing Gmail messages per day related to COVID-19, in addition to more than 240 million COVID-related daily spam messages. Our machine learning models have evolved to understand and filter these threats, and we continue to block more than 99.9 percent of spam, phishing and malware from reaching our users<sup>6</sup>.

Google's TAG has specifically identified over a dozen government-backed attacker groups using COVID-19 themes as lure for phishing and malware attempts—trying to get their targets to click malicious links and download files, including in Australia.

One notable campaign attempted to target personal accounts of U.S. Government employees with phishing lures using American fast food franchises and COVID-19 messaging. Some messages offered free meals and coupons in response to COVID-19, others suggested recipients visit sites disguised as online ordering and delivery options. Once people clicked on the emails, they were presented with phishing pages designed to trick them into providing their Google account credentials.



Location of users targeted by government-backed COVID-19 related attacks

The vast majority of these messages were sent to spam without any user ever seeing them, and we were able to preemptively block the domains using Safe Browsing. We're not aware

<sup>6</sup> Google Threat Analysis Group Blog - [Findings on COVID-19 and online security threats](#)

of any user having their account compromised by this campaign but, as usual, we notify all targeted users with a “government-backed attacker” warning.

We’ve also seen attackers try to trick people into downloading malware by impersonating health organisations, and those organisations themselves are increasingly becoming targets of attacks.

Generally, we’re not seeing an overall rise in phishing attacks by government-backed groups; we are, however, observing a change in tactics. In fact, we saw a slight decrease in overall volumes of phishing activity in March compared to January and February. While it’s not unusual to see some fluctuations in these numbers, it could be that attackers, just like many other organisations, are experiencing productivity lags and issues due to global lockdowns and quarantine efforts.

## Conclusion

Long term success in mitigating disinformation and foreign influence through social media rests on the development of a culture of online safety across society. This includes ongoing collaboration between relevant stakeholders including industry, the technical community and government, as well as efforts at educating users and organisations, from school pupils through to senior citizens and company employees, on how to secure their online presence and to apply critical thinking to the information they see and consume.

Google will continue to invest strongly in threat detection and policy enforcement, and will act decisively to protect our users from disinformation and prevent the abuse of our platforms for government-backed foreign interference operations, phishing and other forms of targeted attacks.

We welcome the Committee’s work in this important policy domain, and look forward to continuing our engagement.

ENDS