MAY 2024

FIRST QUARTER

Adversarial Threat Report

Margarita Franklin, Director, Public Affairs, Security
Dr. Lindsay Hundley, Influence Operations Policy Lead
Mike Torrey, Security Engineer
David Agranovich, Security Policy Director, Threat Disruption
Mike Dvilyanski, Head of Threat Investigations

Meta

TABLE OF CONTENTS

PURPOSE OF THIS REPORT

Our public threat reporting began over six years ago when we first shared our findings about coordinated inauthentic behavior (CIB) by a Russian covert influence operation. Since then, we have expanded our ability to respond to a wider range of adversarial behaviors as global threats have continued to evolve. To provide a more comprehensive view into the risks we tackle, we've also expanded our threat reports to include insights into other threats, as part of our quarterly reporting. In addition, we're also publishing threat indicators to contribute to the security community's efforts to detect and counter malicious activity elsewhere on the internet (see Appendix).

We expect the make-up of these reports to continue to change in response to the changes we see in the threat environment in different areas. This report is not meant to reflect the entirety of our security enforcements, but to share notable trends and investigations to help inform our community's understanding of the evolving threats we see. We welcome ideas from our peers to help make these reports more informative.

For a quantitative view into our enforcement of our Community Standards, including content-based actions we've taken at scale and our broader integrity work, please visit Meta's Transparency Center here: https://transparency.fb.com/data/.

What is Coordinated Inauthentic Behavior or CIB?

We view CIB as coordinated efforts to manipulate public debate for a strategic goal, in which fake accounts are central to the operation. In each case, people coordinate with one another and use fake accounts to mislead others about who they are and what they are doing. When we investigate and remove these operations, we focus on behavior, not content — no matter who's behind them, what they post or whether they're foreign or domestic.

Continuous CIB enforcement: We monitor for efforts to come back by networks we previously removed. Using both automated and manual detection, we continuously remove accounts and Pages connected to networks we took down in the past. See the Doppelganger section for details on our approach to persistent threat threats.

KEY INSIGHTS

In this report, we're sharing threat research into six new covert influence operations that we've disrupted, including from Bangladesh, China, Croatia, Iran, Israel, and a CIB network that targeted Moldova and Madagascar. Many of these cross-internet campaigns were detected and removed early in their audience building efforts. In addition, we're including our latest findings into a long running covert influence operation from Russia, known as Doppelganger. Finally, as part of our mid-year update on the global threat landscape, we're sharing some key insights that stood out to us in our threat research to date:[1]

1. So far, we have not seen novel GenAI-driven tactics that would impede our ability to disrupt the adversarial networks behind them. We've observed instances of: photo and image creation, AI-generated video news readers,[2] and text generation. We have not seen threat actors use photo-realistic AI-generated media of politicians as a broader trend at this time. Here are some examples of what we've seen to date:

- Threat actors continue to use profile photos created using generative adversarial networks (GAN) for their fake accounts, which hasn't impacted our ability to detect inauthentic networks behind them.
- A deceptive network from China shared poster images – likely generated using AI – for a fictitious pro-Sikh activist movement called Operation K.
- An Israel-based CIB network posted likely AI-generated comments under Pages of media organizations and public figures. These comments included links to the operation's websites and were often met with critical responses from authentic users calling them propaganda.

We found and removed many of these campaigns early, before they were able to build audiences among authentic communities. While we continue to monitor and assess the risks associated with evolving new technologies like AI, what we've seen so far shows that our industry's existing defenses, including our focus on behavior (rather than content) in countering adversarial threat activity, already apply and appear effective.

2. We've identified major changes in tactics on Meta's apps by Doppelganger. Focused primarily on weakening international support for Ukraine, this campaign continues to be a "smash-and-grab" effort expending a large amount of resources in the face of a very high detection rate and daily loss of assets. Such persistence is expected for an influence campaign run "at the direction of the Russian Presidential administration" in wartime. Through daily monitoring, detection, blocking and

[1] See our 2024 outlook in Q3'2023 report for prior insights.

[2] As we mentioned in our Q3'2023 report, the China-based influence operation "Spamouflage" was reported by researchers at Graphika to have used AI-generated newsreaders in their videos. These early attempts at using AI-generated videos were quickly identified and exposed. Also, see this report on instances of using GenAI in generating 'news reader' stype propaganda videos by ISIS by Global Network on Extremism and Technology.

exposing Doppelganger's attempts to target our platform since 2022, it's largely ceased to engage in the following tactics *on our apps*, while still actively deploying them elsewhere online:

- No linking to spoofed websites impersonating news media or government agencies;
- No commenting on posts by other people;
- No fictitious brands present on platform (e.g., Reliable Recent News);
- No seeding of links to drive traffic off-platform (in ads, posts, comments, etc.), including via multiple redirects;

While these are significant shifts in on-platform behavior, we know this may change as Doppelganger tries to evolve. Our teams remain vigilant to block new tactics. More details here.

3. While we've seen public discourse ahead of the EU parliamentary elections focus primarily on foreign threats, including from Doppelganger, the majority of the EU-focused inauthentic behavior we've disrupted so far has been domestic in nature. This includes both CIB activity (i.e., Croatia) and simpler inauthentic clusters we removed in recent months in Europe, including in France, Germany, Poland and Italy.[3] Here is what stood out to us:

- Small number of assets within each cluster or network;
- Primarily targeted audiences in their own countries;
- Mostly focused on local elections, rather than the EU parliamentary elections;
- Many were linked to individuals associated with local campaigns or candidates;
- Tactics included inauthentic amplification of authentic accounts or Pages of domestic politicians through likes, shares and comments to make them appear more popular than they were;
- We haven't seen evidence of these clusters gaining much traction among authentic audiences.

On the foreign threats side, the attempts we've seen so far (including Doppelganger and a handful of IB clusters we took down) were primarily focused on undermining support for Ukraine among the EU member states, rather than directly targeting the EU parliamentary elections. More details here.

[3] CIB networks typically involve the use of fake accounts in an identity-based deception, whereas the IB activity is primarily centered around amplifying and increasing the distribution of content with little attempt to obfuscate threat actors' identity from Meta and only superficial attempts to construct a false identity.

01

Bangladesh

We removed 50 accounts on Facebook and 98 Pages for violating our policy against coordinated inauthentic behavior. This network originated in Bangladesh and targeted domestic audiences in that country.

The people behind this activity used fake accounts – some of which were detected and disabled by our automated systems prior to our investigation – to post content and manage Pages. Some of these Pages posed as fictitious new entities and some used names of existing news organizations in Bangladesh. A few Pages used the Bangladesh Nationalist Party (BNP) in their name and posted anti-BNP content. Many of these Pages had a corresponding presence across several platforms, including YouTube, X (formerly Twitter), TikTok and Telegram, in addition to their own websites.

The network posted primarily in Bengali and also in English about news and current events in Bangladesh, including elections, criticism of the BNP, allegations of BNP's corruption and its role in pre-election violence, as well as supportive commentary about the incumbent government, the ruling party and its role in the technological development of Bangladesh.

We found this activity as a result of our internal investigation into spammy inauthentic amplification activity in the region that we removed last year, which led us to uncover a separate coordinated inauthentic behavior network reported here. Although the people behind it attempted to conceal their identity and coordination, our investigation found links to individuals associated with the Awami League party and the Center for Research and Information, a non-profit in Bangladesh.

- *Presence on Facebook and Instagram:* 50 Facebook accounts and 98 Pages.
- *Followers:* About 3.4 million accounts followed one or more of these Pages.
- Ad spend: About $60 in spending for ads on Facebook, paid for mostly in Bangladeshi taka.

02

China

We removed 37 Facebook accounts, 13 Pages, five Groups, and nine accounts on Instagram for violating our policy against coordinated inauthentic behavior. This network originated in China and targeted the global Sikh community, including in Australia, Canada, India, New Zealand, Pakistan, the UK, and Nigeria.

This activity – targeted at multiple services, including ours, Telegram and X (former Twitter) – included several clusters of fake accounts, including one with links to an unattributed CIB network from China targeting India and the Tibet region that we disrupted in early 2023. Some of these clusters amplified one another with most of their engagement coming from their own fake accounts, likely to make this campaign appear more popular than it was. This operation used compromised and fake accounts – some of which were detected and disabled by our automated systems prior to our investigation – to pose as Sikhs, post content and manage Pages and Groups. They appeared to have created a fictitious activist movement called Operation K which called for pro-Sikh protests, including in New Zealand and Australia. We found and removed this activity early, before it was able to build an audience among authentic communities.

They posted primarily in English and Hindi about news and current events, including images likely manipulated by photo editing tools or generated by artificial intelligence, in addition to posts about floods in the Punjab region, the Sikh community worldwide, the Khalistan independence movement, the assassination of Hardeep Singh Nijjar, a pro-Khalistan independence activist in Canada, and criticism of the Indian government.

We found this activity as part of our internal investigation into suspected coordinated inauthentic behavior in the region.

- *Presence on Facebook and Instagram:* 37 Facebook accounts, 13 Pages, five Groups, and nine Instagram accounts.
- *Followers*: About 2,700 accounts followed one or more of these Pages, about 1,300 accounts joined one or more of these Groups, and under 100 accounts followed one or more of these Instagram accounts.

03

Croatia

We removed 104 Facebook accounts, 39 Pages, and seven accounts on Instagram for violating our policy against coordinated inauthentic behavior. This network originated in Croatia and targeted domestic audiences in that country.

The individuals behind this activity used fake accounts — some of which were detected by our automated systems prior to this investigation — to run Pages, create fictitious personas, share and comment on other people's posts. As this network failed to build an audience for its own Pages, it shifted its focus to commenting on news and political Pages in Croatia. It included supportive commentary on the official Pages for the Croatian Democratic Union (the HDZ party) and its politicians and candidates — both national and local. These comments rarely received authentic engagement from other people.

This campaign's own accounts befriended each other, used GAN profile photos, and liked comments by other fake accounts in this network, in an attempt to make them appear more popular than they were. One of the network's fictitious personas posed as an individual with a military background and posted memes about right-wing candidates in the Croatian elections as puppets of Victor Orbán of Hungary and Vladimir Putin of Russia. Another persona posed as an environmentalist and posted about its distrust of the media, criticisms of the mayor of Zagreb, and praise for HDZ.

This network posted primarily in Croatian about issues related to the national elections, economy, successes of the HDZ-led government, as well as critical commentary about opposition figures, the President of Croatia and his Social Democratic Party.

We began looking into this activity after reviewing public reporting about a portion of this network by a Croatian NGO Gong. Notably, this operation moved to change the profile pictures and names of its fake accounts after Gong's report. While the people behind this network attempted to conceal their identity and coordination, we found links to individuals associated with HDZ's youth organization Mladež Hrvatske Demokratske Zajednice (MHDZ).

- *Presence on Facebook and Instagram:* 104 Facebook accounts, 39 Pages, and seven Instagram accounts.
- *Followers*: About 100 accounts followed one or more of these Pages, and about 150 accounts followed one or more of these Instagram accounts.

04

Iran

We removed 22 Facebook accounts, eight Pages, eight Groups, and 23 accounts on Instagram for violating our policy against coordinated inauthentic behavior. This network originated in Iran and targeted Israel.

This network included several separate clusters of activity and used fake accounts to create fictitious personas posing as Israelis in Israel and abroad, manage Groups and Pages and post content. This operation had presence across the internet, including on Telegram, YouTube, X (former Twitter), and TikTok, likely to backstop its fictitious personas so they appear more legitimate and can withstand scrutiny. It showed relatively consistent operational security (OpSec) aimed to conceal their origin. We found and removed these clusters before they were able to build an audience among authentic communities.

The individuals behind this activity posted primarily in Hebrew about news and current events in Israel, including criticism of Hamas and supportive commentary about Israel. One cluster of accounts purported to be right-leaning individuals who posted content in support of Prime Minister Netanyahu prior to the October 7 attack by Hamas. Soon after, they began posting in support of ultra-conservative policies and an Israeli politician, Itamar Ben-Gvir, while criticizing Netanyahu's handling of Israel's response to the attack. Other clusters included fictitious news entities that posted about the Haredi community and ultra-Orthodox demonstrations. A separate subset of accounts posed as young liberal Israeli women and posted about anti-government protests. We saw instances where people called them out as fake. Yet another cluster posed as liberal Israelis posting in support of the LGBTQ+ community.

We began looking into this activity as a result of our internal investigation into suspected coordinated inauthentic behavior in the region. Our analysis benefited from public research by the Israeli Democracy Institute.

- *Presence on Facebook and Instagram:* 22 Facebook accounts, eight Pages, eight Groups, and 23 Instagram accounts.
- *Followers:* About 900 accounts followed one or more of these Pages and about 1,400 accounts joined one or more of these Groups, and about 3,200 followed one or more of these Instagram accounts.

05

Israel

We removed 510 Facebook accounts, 11 Pages, one Group, and 32 accounts on Instagram for violating our policy against coordinated inauthentic behavior. This network originated in Israel and primarily targeted audiences in the United States and Canada.

This cross-internet operation targeted many services, including ours, X (formerly Twitter), and YouTube, and operated several distinctly branded websites focused on the Israel-Hamas war and Middle Eastern politics. We found and removed this network early in its audience building efforts, before they were able to gain engagement among authentic communities.

The individuals behind this activity used fake and compromised accounts, a large portion of which were detected and continuously disabled by our automated systems, before our investigation even began. As these accounts kept going down, the people behind them kept on adding others, likely acquired from account farms, which were also detected and taken down. In addition, this campaign appeared to have purchased inauthentic engagement (i.e. likes and followers) from Vietnam in an attempt to make its content appear more popular than it was.

This network commented on Facebook Pages of international and local media organizations, and political and public figures, including US lawmakers. Their comments included links to the operation's websites and were often met with critical responses from authentic users calling them propaganda. We assess that some of these text comments were likely generated using artificial intelligence. Many of them were not related to the posts they responded to, which is a tactic we've seen in unsophisticated CIB campaigns and spam.

This network's accounts posed as locals in the countries they targeted, including as Jewish students, African Americans and 'concerned' citizens. They posted primarily in English about the Israel-Hamas war, including calls for the release of hostages; praise for Israel's military actions; criticism of campus antisemitism, the United Nations Relief and Works Agency (UNRWA), and Muslims claiming that 'radical Islam' poses a threat to liberal values in Canada.

We began looking into this activity after reviewing public reporting about inauthentic behavior on X (formerly Twitter) by DFRLab at the Atlantic Council, which led us to find corresponding activity on our apps. As our investigation was ongoing, DFRLab published their follow-on research that included a portion of the activity in this report which we confirm to be connected to their original findings on X as part of the same operation. This campaign demonstrated a relative discipline in maintaining OpSec, including by leveraging North American proxy infrastructure to anonymize its activity. While the individuals behind it attempted to conceal their identity and coordination, we

found links to STOIC, a political marketing and business intelligence firm based in Tel Aviv, Israel. It is now banned from our platform. We issued a cease and desist letter to STOIC, demanding that they immediately stop activity that violates Meta's policies.

- *Presence on Facebook and Instagram:* 510 Facebook accounts, 11 Page, one Group, and 32 Instagram accounts.
- *Followers:* About 500 accounts followed one or more of these Pages, less than 100 accounts joined this Group, and about 2,000 accounts followed one or more of these Instagram accounts.

06

Unknown origin

We removed 1,326 Facebook accounts, 80 Pages, one Group and one account on Instagram for violating our policy against coordinated inauthentic behavior. This network targeted audiences in Moldova and Madagascar.

It included several clusters of activity present on multiple internet services, including Facebook, Telegram, Vimeo, and Change[.]org. To obfuscate their origin on our apps, this operation deployed relatively consistent operational security (OpSec), including using browser anonymizers and VPN and proxy infrastructure out of Lithuania and the Transnistria region of Moldova. However, their operational slip-ups led us to uncover and connect these efforts together as part of related deceptive activity.

The individuals behind this operation used compromised and fake accounts, many of which were detected and removed prior to this investigation by our automated systems and investigative teams focused on inauthentic behavior in Moldova. They used these accounts – some of which had GAN profile photos – to pose as locals, manage Pages and Groups, and drive people to off-platform content posted by this network on other services.

This network posted primarily in Romanian, Russian and Malagasy, and also in French about news and current events in Moldova and Madagascar. In Moldova, they posted about the Gagauzia region, public figures including Ilan Shor, a sanctioned Moldovan politician, and Marina Tauber, a member of the Șor Party, in addition to criticizing the incumbent government and its efforts towards EU integration, including with parody videos about the current President. At times, they posted about non-political topics like soccer and dating, likely to appear more authentic.

In Madagascar, they focused primarily on promoting the politician Siteny Randrianasoloniaiko and criticizing the incumbent President Andry Nirina Rajoelina. These attempts to amplify events related to the opposing presidential candidate had no substantive engagement from authentic audiences. This operation also created a petition on Change[.]org claiming that the current president is illegitimate due to his French citizenship.

We found this activity as a result of our internal investigation into suspected coordinated inauthentic behavior in Eastern Europe.

- *Presence on Facebook and Instagram:* 1,326 Facebook accounts, 80 Pages, one Group and one Instagram account.

- *Followers:* About 20,000 accounts followed one or more of these Pages, under 200 accounts joined this Group, and about 10 accounts followed this Instagram account.
- *Ads*: Around $42,000 in spending for ads on Facebook, paid for mostly in US dollars and euros.

07

Russia

‘DOPPELGANGER’S ATTEMPTS TO STAY AFLOAT ACROSS THE INTERNET

As part of our ongoing transparency reporting on Doppelganger, a cross-internet influence operation from Russia, we’re sharing our 7th update in 20 months that includes our latest research findings into this malicious activity. It includes: a major shift in this operation’s tactics on our platform and its latest attempts at evading detection, in addition to publishing another 600+ threat indicators to our industry’s largest repository of 5,000+ indicators related to this threat actor so that our peers and researchers can investigate and take action as appropriate.

What is Doppelganger? Nearly two years ago, we were the first technology company to publicly report on Doppelganger, an operation centered around a large network of websites spoofing legitimate news outlets. The EU Disinfo Lab and the Digital Forensic Research Lab published open source research at the same time. In December 2022, we were first to publicly attribute it to two companies in Russia who were sanctioned by the EU in 2023 and by the US Treasury Department in 2024.

PERSISTENCE & ONGOING ADVERSARIAL ADAPTATION

An APT of influence operations

While persistence is common among influence operations, Doppelganger has taken it to a new level over the last 20 months, while remaining crude and largely ineffective in building authentic audiences on social media. Any research into its activity should treat this cross-internet campaign as **an advanced persistent threat**, rather than a short-lived operation, if the goal is to properly analyze it and develop counter-measures.

Quantity over quality

Since our last report in February, Doppelganger has continued its “smash-and-grab” efforts, expending a large amount of resources – even when it leads to a very high detection rate and daily loss of assets. However, the operators’ efforts to evade detection in response to our aggressive blocking of their recidivist attempts have resulted in degrading the quality of the overall operation (see details below). This is consistent with assessments by other

researchers throughout Doppelganger's lifespan.[4]

Our expectations

Just like for any security team in our industry or in government, it would be unrealistic to expect advanced persistent threat actors to cease their activity after their online accounts are taken down by any one platform or when they are sanctioned by any one government. This is particularly true for a persistent influence campaign run by commercial companies "at the direction of the Russian Presidential Administration" in wartime. For-hire groups are paid to keep at it for as long as their clients' operational objectives remain – in this case, to undermine the international community's support for Ukraine.

Perception hacking risks

We also know that persistent and motivated operations like Doppelganger – in response to constant detection and blocking of their accounts and websites – may leverage their own notoriety to create the perception that they are more impactful than they are. They may do it to justify their budgets to their clients or to undermine people's trust in the information environment. Overstating their impact can have the effect of eroding the very concept of "facts" and sowing distrust in electoral outcomes and public institutions.

We remain focused on making these attempts more costly and less effective, including by routinely sharing information about what we see.

OUR APPROACH TO COUNTERING DOPPELGANGER

Without going into details that may tip off malicious groups on how to evade detection, our teams are engaged in *daily efforts* to find and block Doppelganger's attempts to acquire new accounts, run ads, and share links to its websites and redirect domains, before these are ever shared on our apps.

---

[4] ISD, Pro-Kremlin Network Impersonates Legitimate Websites and Floods Social Media with Lies, September 29, 2022: *"The actors behind the operation have expended little effort on making the content look authentic and seem to have concentrated on distribution volume instead. Most of the posts that ISD analysed on Facebook and Twitter have received little to no authentic interaction."*

Euronews, 'Doppelganger': How France exposed a massive Russian disinformation campaign, June 2023: *"According to the French government, so far, this vast disinformation campaign has been mostly unsuccessful in garnering clicks and drawing attention."*

Recorded Future, Russia-linked 'Doppelgänger' social media operation rolls on, report says, *December 2023: "Despite the campaign's high volume of CIB, we did not identify any significant engagement from authentic social media users with any of the articles. Viewership and other engagement metrics (reshares, likes, and replies) were negligible across the network."*

Over time, we've seen that forcing them to adversarially adapt as we continue to improve our defenses degrades the quality of their operation overall. For example, at the start, Doppelganger used relatively convincing typo-squatting domains. After we blocked these domains from being shared on our apps, the operators had to begin using redirects, including with random and unrelated strings in urls, making these attempts more obviously fake. *(See more examples of their latest text obfuscation tactics in ads below)*

At this stage, and after nearly two years of investigating this operation, it looks to be as persistent and voluminous in its attempts as spammers are in targeting people online with knockoff merchandise: constantly shifting key words, spelling, off-platform links, and images, and churning through many burner accounts and Pages to only leave a single comment or run a single ad before we block them.

Our ongoing work to counter Doppelganger is three-fold:

1. Pre-empt or disrupt the immediate activity and force the operators to rebuild and re-tool over and over again;
2. Feed new detection signals into automated enforcement systems as we learn from each iteration in tactics so we can enforce at a global scale;
3. Share information with industry peers and relevant governments to help everyone raise defenses. Since 2022, we've shared our insights with many companies targeted by Doppelganger and also government officials in countries around the world, including the US, France, Germany, Ukraine, UK, Poland, Latvia, Ireland, Australia, European Commission and EU External Action Service.

LATEST INSIGHTS

1. Major shift in tactics on our platform, unmatched by activity on other services

As of this reporting, our research shows the following changes in Doppelganger's use of its tactics on our apps, while it's still actively deploying them elsewhere online:

- No linking to spoofed domains impersonating news media or government agencies;
- No commenting on posts by other people and organizations;
- No fictitious brands present on our apps (e.g., Reliable Recent News, etc.);
- No seeding of links to drive traffic to off-platform domains (e.g., via ads, posts, comments, etc.);

Given how central the websites have been to this campaign from the start, we've heavily focused on reducing Doppelganger's ability to seed these links directly or through redirect urls. Since our last threat report in February, we have not seen Doppelganger achieve much success getting

people to visit their websites by seeding links on our apps, despite all its efforts. As of April, the operators have stopped attempting to share links altogether.

While these are significant shifts in on-platform behavior, we know this may change as Doppelganger keeps trying to find ways around detection, and our teams remain vigilant to respond.

2. Attempts at text obfuscation to avoid detection

Among the most recent adversarial adaptations, which we have now incorporated into our detection systems, we've seen the operators use text obfuscation tactics to avoid detection. This may include adding combinations of whitespace and extra punctuation to break up words (such as "U. kr. ai. n. e" instead of "Ukraine"), or even hidden space characters so that every word in the ad is broken up with a zero-width space character in unicode in between every letter. Breaking up words in these ways makes these ads barely legible and nonsensical, raising questions about the actual intent behind these efforts, beyond just checking the box that the ad attempt was made. In fact, authentic users commented on these ads, calling them out as Russian trolls, propaganda and bots.

Images: examples of ads using text obfuscation techniques

Overall, the majority of Doppelganger ad attempts have been overtly pro-Russian, putting no effort in hiding its focus on promoting Russia and criticizing Ukraine in the context of the ongoing war. Here are its five top-performing ads globally, which were aimed at Algeria, Pakistan, India, Poland and Venezuela.[5]

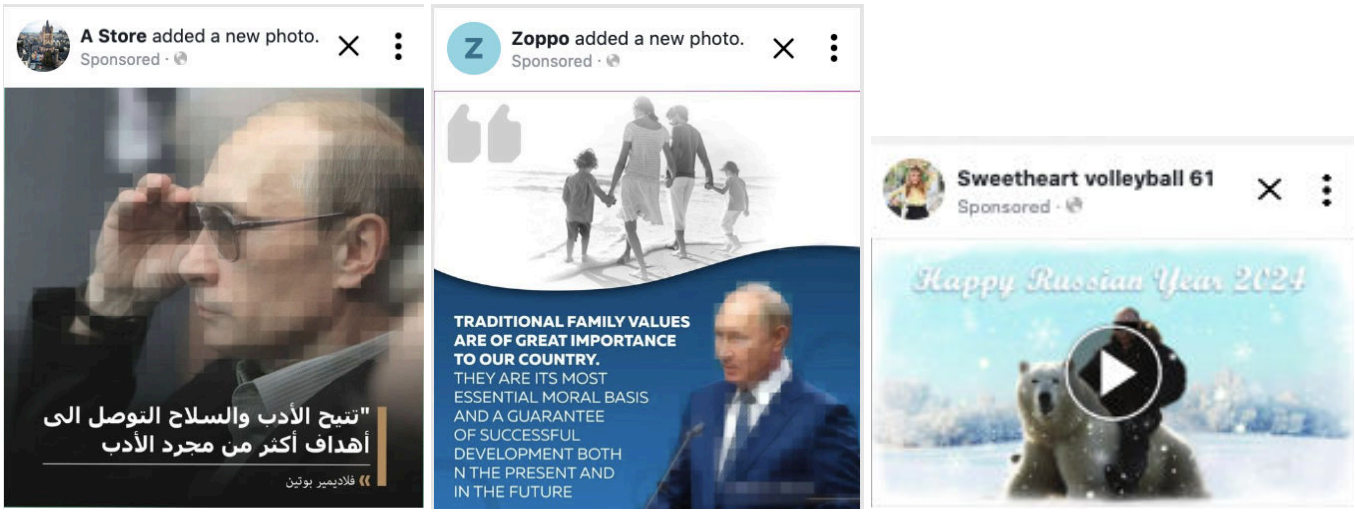

Images

First: an ad targeting audiences in Algeria; Second: an ad targeting audiences in Pakistan; Third: an ad targeting audiences in India

Images

Fourth: an ad targeting audiences in Poland; Fifth: an ad targeting audiences in Venezuela

[5] While the vast majority of activity by Doppelganger over two years has focused on Ukraine, Germany, France, US, and Israel and to a lesser extent on Italy and Poland, we have also observed occasional brief targeting of other countries in the world.

RECOMMENDATIONS FOR STRONGER INDUSTRY RESPONSE

Like most of the threats we see, Doppelganger does not limit its targeting to one service. Since the start, it operated across many platforms, including Facebook, Instagram, Telegram, X, YouTube, TikTok, and even LiveJournal, petition websites, and small blogging services, in addition to the thousands of websites and redirect domains it controls.

While many of our peers, researchers and governments are monitoring for malicious activity by this persistent campaign and others, it would benefit us all to have a fuller shared picture of the global activity by these operators across the internet. We have provided routine updates on what we see on our apps, but we only have a limited view into these malicious efforts across the internet. In fact, the behavior we see can sometimes appear somewhat puzzling from only our vantage point. For example, while Doppelganger has abandoned sharing links to its vast web of websites and domains on our apps, they continue to frequently update these websites with new content, suggesting the websites may still be receiving traffic from elsewhere.

Based on our threat research over nearly two years, here are a few recommendations for how the defender community can continue expanding transparency and information sharing to enable better defenses against persistent threats like Doppelganger:

- **Technology companies:** In addition to sharing our own threat research, platforms should consider making it easier for open-source researchers to map their findings across different services. For example, they can make political ad libraries searchable for key terms used by Doppelganger and ensure that researchers can see content from political ads even after they have been removed as violating.
- **Domain registrars and hosting providers:** Doppelganger's off-platform websites continue to survive for months if not years after exposure by the security community. In addition to the recommendations we made last year to better tackle domain registration abuse, domain registrars and hosting providers should also consider sharing information about new domains that a known malicious operation creates. Because these groups have information about newly registered domains early, proactive sharing can help stymie threat actors before these domains gain traction on the internet.
- **Governments:** Sharing information between tech companies, governments and law enforcement can be critical in disrupting malicious foreign campaigns early, particularly in cases when operators coordinate their efforts outside of any one platform. While we have shared information about Doppelganger's targeting with various government officials and law enforcement in a number of countries, it would be beneficial to understand – when possible – how campaigns like these run their technical and financial infrastructure to enable their global operations. Such insights can help researchers and companies to identify opportunities to disrupt malicious activity more comprehensively.

Over the years, we have seen first-hand how transparency and information sharing can be a force multiplier that enables follow-on threat research and disruptions, raising the cost of running these operations while making them less and less effective across the board.

Appendix: Threat indicators

The following section details unique threat indicators that we assess to be associated with the malicious networks we disrupted and described in this report. To help the broader research community to study and protect people across different internet services, we've collated and organized these indicators according to the Online Operations Kill Chain framework, which we use to analyze many sorts of malicious online operations, identify the earliest opportunities to disrupt them, and share information across investigative teams. The kill chain describes the sequence of steps that threat actors go through to establish a presence across the internet, disguise their operations, engage with potential audiences, and respond to takedowns.

We're sharing these threat indicators to enable further research by the open-source community into any related activity across the web (GitHub). This section includes the latest threat indicators and is not meant to provide a full cross-internet, historic view into these operations. It's important to note that, in our assessment, the mere sharing of these operations' links or engaging with them by online users would be insufficient to attribute accounts to a given campaign without corroborating evidence.

BANGLADESH-BASED CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	50 accounts
Acquiring Facebook Pages	98 Pages
Acquiring domains to support influence operations	dhakatv[.]net
	bnppara[.]com
	bdpolitico[.]com
	bdpolitico[.]org
	banglanewsbank[.]com
	bdperspectives[.]com

	bangladeshtimes360[.]com
	londonbanglanews[.]com
	bdanalytica[.]com
	bnpnews[.]net
	hyperjoshim[.]com
	bdsarcasm[.]net
	khambastar[.]net
	news360[.]com
	newsbangladesh[.]org[.]bd
	todaybd24[.]com
Acquiring X accounts	https://twitter[.]com/bd_politico
	twitter[.]com/tonmoybuet
Acquiring Telegram channels	https://t[.]me/bnppara
Acquiring YouTube channels	https://www[.]youtube[.]com/@DhakaTelevision
	https://www[.]youtube[.]com/@bnppara
	https://www[.]youtube[.]com/@bdpolitico1527
	https://www[.]youtube[.]com/@hyperjoshim
	https://www[.]youtube[.]com/channel/UCHa42HrKBenuE2BdmOHqlfA
	youtube[.]com/@bnpnama434
	youtube[.]com/@banglapolitix
Acquiring TikTok accounts	https://tiktok[.]com/@bnppara
	https://www.tiktok[.]com/@bnpnewsbd
Disguising assets	
Adopting a visual disguise	Copying profile pictures

Posing as non-existent entity	The network posed as fictitious news outlets: Dhaka Television London Bangla News BNP Para Bangla Politix Bangladesh Perspectives BDAnalytica
Impersonating real institutions	The network used names of existing news organizations in Bangladesh
Backstopping	Pages had presence across Youtube, X (formerly Twitter), TikTok, Telegram, and their own websites
Gathering Information	
Monitoring specific events	The network was monitoring breaking news, and wrote about electoral protests soon after it occurred
Coordinating and planning	
Coordinating shift patterns	The network was working to a regular shift pattern, posting between 0700 and 2100 GMT, with a peak between 1300-1800 GMT and notably fewer posts on Fridays
Evading detection	
Privacy protection	Domains created by the network were registered under privacy protection
Indiscriminate engagement	
Posting on websites	Posted articles on websites controlled by the operation

Targeted engagement	
Running ads	*Advertising to promote content, goods and services : About $60 in spending for ads on Facebook, paid for mostly in Bangladeshi taka*
Posting about individuals or institutions	*Posted negative commentary about the BNP and BNP politicians*
	Posted positive commentary about the Awami League and the PM

CHINA-BASED CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	37 accounts
Acquiring Facebook Groups	5 Groups
Acquiring Facebook Pages	13 Pages
Acquiring Instagram accounts	9 accounts
Acquiring X accounts	twitter[.]com/olivn388
	twitter[.]com/AishaHadi19
	twitter[.]com/SatbirS06914749
	twitter[.]com/s8175593
	twitter[.]com/Asimov_47
Acquiring Telegram channels	http://t[.]me/+ugoxkjuomho4nmy1
Acquiring and repurposing accounts	The network used compromised accounts
Disguising assets	
Posing as non-existent person	Posing as Sikh Activists
Gathering Information	
Monitoring specific events	The network posted about floods in the Punjab region, the Sikh community worldwide, the Khalistan independence movement, the assassination of Hardeep Singh Nijjar, a pro-Khalistan independence activist in Canada

Indiscriminate engagement	
Amplifying with likely fake accounts on social media	The network amplified one another with fake accounts to give the impression the campaign was more popular than it was.

CROATIA-BASED CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	104 accounts
Acquiring Facebook Pages	39 Pages
Acquiring Instagram accounts	7 Instagram accounts
Disguising assets	
Adopting visual disguise	Copying profile photos from online sources
	Using profile photos likely generated using artificial intelligence such as Generative Adversarial Networks (GAN)
Posing as non-existent person	Posing as military personnel
	Posing as environmentalist
Evading detection	
Camouflaging content	Many accounts posted spammy photos or videos of scenery, food, etc., likely to appear more authentic
	Some accounts made 'personal' comments alongside the links they shared, likely to appear more individual
Obfuscating infrastructure	Routing traffic through proxy infrastructure
Indiscriminate engagement	
Amplifying with likely fake accounts on social media	Sharing on Facebook and liking comments made by other fake accounts

Targeted engagement	
Running ads	About 20 USD in spending for ads on Facebook, paid for in US Dollars
Engaging with users outside the operation	About 100 accounts followed one or more of these Pages
	About 150 accounts followed one or more of these Instagram accounts
Engaging with specific audience	Posting into Groups focused on politics
	Commenting on news and political Pages in Croatia
Directing online traffic	Directing audience towards off-platform websites
Posting about individuals or institutions	Commenting positively on the official Pages for the Croatian Democratic Union or the HDZ party, its politicians at both national and local levels
	The network commented critically about opposition figures, the President of Croatia and his Social Democratic Party
Enabling longevity	
Replacing Infrastructure	The network changed various fake accounts' names and profile pictures after public reporting exposing some of this activity

IRAN-BASED CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	22 accounts
Acquiring Facebook Groups	8 Groups
Acquiring Facebook Pages	8 Pages
Acquiring Instagram accounts	23 accounts
Acquiring domains to support influence operations	http://jgtm[.]org/
	tikvaodesa[.]com
Acquiring X accounts	https://twitter[.]com/changebeliever
	https://twitter[.]com/SecularIsraelis
	https://twitter[.]com/JgtmOrg
	https://twitter[.]com/jgtm_offical
	https://twitter[.]com/KodeshNews
Acquiring TikTok accounts	tiktok[.]com/@secularisrael
	https://www.tiktok[.]com/@jewishfist
	https://www.tiktok[.]com/@israel_breaking_news
Acquiring Telegram channels	https://t[.]me/IsraeltheSecond
	https://t[.]me/Israel_BreakingNews

	https://t[.]me/newsrespond
	https://t[.]me/respondnewsil
	https://t[.]me/newsnowil
	https://t[.]me/KodeshNews
	https://t[.]me/jgtm_offical
	https://t[.]me/JewishFist
Disguising assets	
Posing as non-existent person	Posing as anti-war activists
	Posing as a local in support of anti-government protests
Posing as non-existent institution	Posing as a fictitious news channel posting about the Haredi community and ultra-Orthadox demonstrations
Backstopping	Creating fictitious personas on other internet services including Telegram, Youtube, X (former Twitter), and Tiktok
Targeted engagement	
Posting about individuals or institutions	Posting content in support of Prime Minister Netanyahu
	Posting content in support of Israeli politician, Itamar Ben-Gvir
	Posting criticism of Hamas

ISRAEL-BASED CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	510 accounts
Acquiring Facebook Groups	1 Group
Acquiring Facebook Pages	11 Pages
Acquiring Instagram accounts	32 accounts
Acquiring YouTube channels	https://www[.]youtube[.]com/@UC4Canada
Acquiring domains to support influence operations	nonagenda[.]com
	ufnews[.]io
	uc4canada[.]com/
	the-good-samaritan[.]com
Disguising assets	
Posing as non-existent person	Posing as Jewish students, African Americans and 'concerned' citizens
Posing as non-existent institution	Creating fictitious news outlets: Nonagenda; Ufnews; uc4canada
Evading detection	
Obfuscating infrastructure	Leveraging North American proxy infrastructure to anonymize activity
Indiscriminate engagement	
Amplifying with fake accounts on	The campaign purchased inauthentic engagement (i.e. likes and followers) from Vietnam in an attempt to make its content appear

Facebook	more popular than it was
Targeted engagement	
Engaging with users outside the operation	About 500 accounts followed one or more of these Pages
	Less than 100 accounts joined the Group
	About 2,000 accounts followed one or more of these Instagram accounts.
Engaging with specific audience	The network commented on Facebook Pages of media organizations – both international and local, and political and public figures, including US lawmakers
Directing online traffic	The network comments linked to the operation's websites
Posting about individuals or institutions	The network posted supportive commentary about Israel's military actions
	The network posted criticisms of Muslims, and the United Nations Relief and Works Agency (UNRWA)
Enabling longevity	
Replacing Infrastructure	Replacing disabled accounts with new fake accounts

UNKNOWN ORIGIN CIB NETWORK

Tactic	Threat indicator
Acquiring assets	
Acquiring Facebook accounts	1326 accounts
Acquiring Facebook Groups	1 Group
Acquiring Facebook Pages	80 Pages
Acquiring Instagram accounts	1 accounts
Acquiring Telegram channels	http://t[.]me/sanduofficial
Acquiring accounts on online forums	http://change[.]org/p/tsy-afaka-mirotsaka-ho-fidiana-fo-filoham-pirenena-intsony-ny-rajoelina/nftexp/fht-37311859-en-gb/cv_287912/1315562991

CONTINUOUS ENFORCEMENT: LATEST THREAT INDICATORS RELATED TO RECIDIVIST ATTEMPTS

We monitor for, and enforce against, efforts to come back by networks we previously removed. Some of these networks may attempt to create new off-platform entities, such as websites or social media accounts, as part of their recidivist activity.

We're sharing some of these novel threat indicators related to recidivism attempts to enable further research by the open-source community into any related activity across the internet. It's important to note that, in our assessment, the mere sharing by online users of these operations' links or engaging with them would be insufficient to attribute these accounts to a given campaign without corroborating evidence.

DOPPELGANGER: LATEST BRANDS & SPOOFED DOMAINS

This section includes the latest domains spoofing news websites that we've identified as part of the Doppelganger campaign as of May 20, 2024.

In addition to these domains, we've identified hundreds more that the campaign uses to redirect people to its spoofed and branded domains. We've updated our full list of threat indicators linked to Doppelganger on GitHub in a machine-readable format.

Domains spoofing news sites

Domain	Registration date	Country likely targeted
leparisien[.]top	2024-03-05	France
lepoint[.]wf	2024-03-31	France
welt[.]pm	2024-01-11	Germany
lastampa[.]in	2024-04-05	Italy
repubblica[.]in	2024-04-05	Italy
polskieradio[.]icu	2024-04-06	Poland