

MARCH 2023

Phase-based Tactical Analysis of Online Operations

Ben Nimmo and Eric Hutchins

Phase-based Tactical Analysis of Online Operations

Ben Nimmo and Eric Hutchins

Carnegie's Partnership for Countering Influence Operations is grateful for funding provided by the William and Flora Hewlett Foundation, Craig Newmark Philanthropies, the John S. and James L. Knight Foundation, Microsoft, Meta, Google, Twitter, and WhatsApp. The PCIO is wholly and solely responsible for the contents of its products, written or otherwise. We welcome conversations with new donors. All donations are subject to Carnegie's donor policy review. We do not allow donors prior approval of drafts, influence on selection of project participants, or any influence over the findings and recommendations of work they may support.

© 2023 Carnegie Endowment for International Peace. All rights reserved.

Carnegie does not take institutional positions on public policy issues; the views represented herein are those of the author(s) and do not necessarily reflect the views of Carnegie, its staff, or its trustees.

No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from the Carnegie Endowment for International Peace. Please direct inquiries to:

Carnegie Endowment for International Peace
Publications Department
1779 Massachusetts Avenue NW
Washington, DC 20036
P: + 1 202 483 7600
F: + 1 202 483 1840
CarnegieEndowment.org

This publication can be downloaded at no cost at CarnegieEndowment.org.


Contents

Summary	1
Introduction	3
Using the Online Operations Kill Chain	5
Principles of the Online Operations Kill Chain	8
The Online Operations Kill Chain	10
Appendix: Case Studies	19
About the Authors	27
Notes	29
Carnegie Endowment for International Peace	35

Summary

The online threatscape in 2023 is characterized by an unprecedented variety of actors, types of operation, and threat response teams. Threat actors range from intelligence agencies and troll farms to child-abuse networks. Abuses range from hacking to scams, election interference to harassment. Responders include platform trust-and-safety teams, government agencies, open-source researchers, and others. As yet, these responding entities lack a shared model to analyze, describe, compare, and disrupt the tactics of malicious online operations. Yet the nature of online activity—assuming the targets are human—is such that there are significant commonalities between these abuse types: widely different actors may follow the same chain of steps. By conducting a phase-based analysis of different violations, it is possible to isolate the links in the chain within a unified model, where breaking any single link can disrupt at least part of the operation, and breaking many links—“completing the kill chain”—can disrupt it comprehensively. Using this model will allow investigators to analyze individual operations and identify the earliest moments at which they can be detected and disrupted. It will also enable them to compare multiple operations across a far wider range of threats than has been possible so far, to identify common patterns and weaknesses in the operation. Finally, it will allow different investigative teams across industry, civil society, and government to share and compare their insights into operations and threat actors according to a common taxonomy, giving each a better understanding of each threat and a better chance of detecting and disrupting it.

Links in the Kill Chain



Introduction

Governments,¹ nonprofit organizations,² commercial companies,³ academic institutions,⁴ and social media platforms⁵ have all invested heavily in setting up teams to tackle some of the abuses within the online environment. In parallel, countries and international institutions have begun work to define and regulate the online space, with initiatives such as the UK's Online Safety Bill (formerly Online Harms Bill)⁶ and the EU's revised Code of Practice on Disinformation⁷ and Digital Services Act.⁸

Underpinning these efforts, the research community has conducted foundational work to define and describe the taxonomy of different threats. The cyber espionage community has led the way with the seminal Intrusion Kill Chain,⁹ the Unified Kill Chain,¹⁰ the MITRE ATT&CK framework,¹¹ the Diamond Model of intrusion analysis,¹² and the Pyramid of Pain approach to prioritizing detection indicators.¹³ In the field of influence operations, a number of experts and organizations have proposed kill chains, including Bruce Schneier,¹⁴ Clint Watts,¹⁵ the Center for Security and Emerging Technology at Georgetown University,¹⁶ and the Credibility Coalition Misinfosec Working Group (AMITT and DISARM frameworks).¹⁷ The Digital Shadows Photon Research Team has proposed a kill chain for account takeovers;¹⁸ Optiv has a cyber fraud kill chain.¹⁹ While many of these reference the Intrusion Kill Chain as their inspiration, each is tailored to a specific violation type, such as hacking, influence operations, or fraud.

These models vary in audience and focus. Some are designed for use by specific defenders—for example, the Intrusion Kill Chain, which offers network defenders an intelligence-based framework to disrupt computer exploitation and attack, or Watts's Social Media Kill Chain,

which proposes a model for social media platforms to detect and understand influence operations. Others are broader, such as Schneier's Influence Operations Kill Chain, which recommends countermeasures against influence operations for tech platforms, intelligence agencies, the media, and educators, among others. Some models focus on threat actors' tactics (AMITT: "Create fake Social Media Profiles / Pages / Groups"), while others focus on their overall strategies (Schneier: "find the cracks in the fabric of society"). All have enriched the public debate around online operations and our understanding of the threatscape.

However, two key gaps remain. First, public debate is hampered by the lack of a common taxonomy and vocabulary to analyze, describe, and compare different types of online operations.²⁰ One problem can have many names: for example, within the space of online political interference, different frameworks refer to "disinformation,"²¹ "information operations,"²² "misinformation incidents,"²³ "malinformation,"²⁴ and "influence operations"—terms which may have distinct meanings but are often used interchangeably.²⁵ Simultaneously, one word can have many meanings: the term "exploitation" covers both executing unauthorized code on a victim's system²⁶ and amplifying an influence campaign with bots, trolls, and "useful idiots."²⁷

Second, each model is designed primarily to analyze a single threat activity, be it hacking, influence operations, spam, or fraud. But online operations are amorphous and do not always fit neatly into a single violation type. For example, the operation known as Ghostwriter²⁸ and an unrelated operation from Azerbaijan²⁹ that Meta disrupted both combined hacking and online disinformation. In 2016, Russian military intelligence famously combined hacking, social media activity, planting of articles by fake personas on mainstream media outlets, and weaponized leaking via a third party.³⁰ Analyzing any of these operations through one threat-specific framework carries the risks of missing other important segments of their activity, underenforcing, and reinforcing siloed approaches to tackling different forms of online abuse.

We have designed the Online Operations Kill Chain to fill these gaps by providing an analytic framework that is designed to be applied to a wide range of online operations—especially those in which the targets are human.³¹ These include, but are not limited to, cyber attacks, influence operations, online fraud, human trafficking, and terrorist recruitment. It is our hope that a common framework for investigators across platforms, in the open-source community, and within democratic institutions will enable more effective collaboration to analyze, describe, compare, and disrupt online operations.

Using the Online Operations Kill Chain

The basis of our approach is that, despite their many differences, online operations still have meaningful commonalities. At the most fundamental level—at risk of sounding simplistic—any online operation has to be able to get online. That likely means, at the very least, acquiring an IP address and (depending on the platform) probably an email address or mobile phone number for verification purposes. If the operation runs a website, it will need hosting, administrators, and a content creation platform. If active on social media, it must be able to acquire or create accounts. It will likely try to evade detection by platforms or users by adopting technical and visual disguises, such as stealing a profile picture or obfuscating a piece of code to get past antivirus scanners.³² All these requirements hold true across threat areas, whether the operation is aimed at espionage or election interference, sex trafficking or selling fake Ray-Bans.

The Online Operations Kill Chain builds on those commonalities to propose a unified phase-based framework to analyze many types of operations. It covers the full range of abuses that the threat intelligence teams at Meta routinely tackle, from cyber espionage and influence operations to scams. It is designed to cover multifaceted operations such as Ghostwriter or the Russian military's hack-and-leak operations, as well as simpler ones. Despite this wide coverage, it focuses on identifying the threat actor's specific tactical, technical, and procedural activities.

Analysts and investigators can use the kill chain on three levels, whether they work at a tech platform, an open-source institution, or a government body. First, they can apply it to a single operation and use it to sequence that operation's activity, finding the combination of tactics, techniques, and procedures (TTPs) that would allow for the earliest detection and disruption.³³

As a hypothetical example, if investigators identify that an influence operation is using a particular niche email domain to set up fake social media accounts (kill chain phase: acquiring assets); disguising them with profile pictures generated using generative adversarial networks, or GANs, such as StyleGAN 2 (disguising assets phase); and then using those accounts to spam links to state media websites (indiscriminate engagement phase), then they can prioritize finding ways to detect the combination of email provider and GAN profile picture, which potentially could help in disrupting further fake accounts *before* they post.

Second, they can use the kill chain to compare multiple operations. This can allow them to analyze commonalities between two operations of the same type (such as two harassment campaigns) or between operations of different types (such as a harassment campaign versus a scam), or even to analyze tactical changes in an individual, long-running operation by a particular threat actor by comparing its behavior at different times.³⁴ This, in turn, can provide the necessary data to prioritize countermeasures that could be applied to multiple operations at the same time.

To continue the above hypothetical example, the investigative team could check the kill chain records of other operations to see if the use of either that particular niche domain or StyleGAN 2 images is a recurring pattern. If they find that StyleGAN 2 images have been used by cyber espionage, harassment, and spam networks but the niche email domain has not, they can prioritize finding ways to detect the images, which could enable them to identify many types of operations at an early stage.

Third, and within the limits of privacy regulation, research teams across different disciplines can use the kill chain to share and compare their findings on different operations. Since each investigative team is likely to see different facets of the operation, they can collectively build up a better understanding than any one team could alone.

To extend our hypothetical example further, let us assume that the investigative team shares its kill chain analysis of the initial operation with its peers among tech platforms, law-enforcement institutions, and the open-source community. By pooling their respective insights according to the kill chain's common framework, this community could identify not only the use of that particular email domain and StyleGAN 2 pictures but also other distinguishing features, such as IP addresses; fictitious personas across social media, blogging, and media platforms; and malware. All of these could then be fed back into each team's understanding of the overall operation, possibly empowering more precise and earlier detection.

This approach would make defenses more resilient by enabling investigators on different teams to “complete the kill chain”: identify multiple points at which an operation could be detected and disrupted. It would also increase resilience by allowing teams who specialize in very different areas—for example, scams, harassment of human-rights defenders, and election interference—to compare the operations they see, identify the most common TTPs, and prioritize them for countermeasure development.

Internal Versus External Use

The kill chain is both an analysis tool for investigators and a vehicle to structure communication. It is designed for use within and between platforms, open-source researchers, and governments.

Within institutions, especially platforms, it allows investigative teams to record the TTPs of different operations according to a unified taxonomy and to identify detection leads and points in the chain where the operation can be disrupted. Indicators for internal sharing can be exceptionally granular, including, for example, the combination of IP address, email domain, malware type, and posting pattern that characterizes the malicious operation. Iterative observations can be made to track an operation's changes over time.

Between institutions, the kill chain allows different teams to describe the operations they have uncovered according to a unified taxonomy and to identify the weak points in the chain and the partners who could break those links.³⁵ Given the restrictions of privacy-protection and information-sharing arrangements, such communication will likely be less granular or comprehensive. It could, however, mean sharing technical indicators such as IP addresses between industry peers and sharing behavioral indicators with the public, such as the distinctive pairs of URLs posted by the Chinese influence operation that Meta disrupted in late 2021.³⁶

We designed the kill chain to be used by the open-source community as well as platforms (see box 1). We have experienced firsthand how much information open-source researchers can uncover.³⁷ The Online Operations Kill Chain is designed to enable them to structure and share their own research in a standardized way. For example, an open-source unit that identifies the websites, social media assets, naming conventions, and posting patterns of an operation based on publicly available information—all of which elements have featured in open-source discoveries before—can use the kill chain to set these out in sequence for the benefit of the public and platforms.

BOX 1: Seeing and Sharing

There are significant differences in the sorts of indicators that different members of the defender community can be expected to see. Social media companies and tech providers are more likely to have consistent insights into the infrastructure that underpins different operations on their platforms; open-source researchers are more likely to have consistent insights into online operations' behavior across many platforms. Moreover, different operations leave very different footprints: a complex, public-facing influence operation will spread across far more surfaces than a spearphishing campaign.

However, these differences should not be overstated: open-source techniques can, under some circumstances, expose many details of an operation's infrastructure. Moreover, the technical indicators that each platform or provider sees may also vary markedly. No one investigative team—whether platform, government, or open-source—has a monopoly on insights into online operations.

This is why we believe that responsible sharing is crucial to enable a comprehensive response to any given abusive operation. What seems a tangential insight to one team may be the precise detail that another team needs to break open the case, so the best way to defend against online operations is for each member of the defender community to share what information they can, together with their contextual assessment of how each indicator fits into the overall operation.

Principles of the Online Operations Kill Chain

We have built the Online Operations Kill Chain according to the following principles:

Observation-based: The Online Operations Kill Chain is restricted to TTPs that can be directly observed, such as an operation's use of internet infrastructure, or demonstrated with high confidence, such as an operation's use of an encrypted messaging app if an asset links to that app in its bio. It is not designed to track activity that can only be hypothesized, such as an operation's strategic goal.

Tactical: The kill chain is designed for tactical analysis of online operations. It is not designed to analyze larger phenomena, such as organic social movements, or measure very large-scale vulnerabilities, such as the overall health of a body politic.

Platform-agnostic: We have designed the kill chain to apply to all kinds of platforms—not only social media, but websites and email providers, for example. Some TTPs include real-world activity, such as setting up shell companies or physical offices, or co-opting influencers, journalists, and others to carry out influence activities, as some troll farms are known to have done.³⁸ The precise activity will vary from one surface to another, but the links in the kill chain are constant.

Optimized for human-on-human operations: We have optimized the Online Operations Kill Chain to describe operations in which the source and target are human—for example, an espionage team trying to socially engineer a diplomat, an influence operation trying to co-opt a journalist, or a network sharing child sexual abuse material. The kill chain can be applied to machine-on-machine attacks, but it is not primarily designed with them in mind.

One or many platforms: We have designed the kill chain to be applicable to both single-platform and multiplatform operations. A number of techniques and procedures explicitly reference cross-platform activity, such as backstopping personas by maintaining the same fake identity on multiple social media and using each platform to boost the credibility of the others, running phishing websites, posting content from one platform to another, and switching conversations from direct messages to emails.

Modular: The kill chain reflects the possible phases of an operation, but not every operator goes through every phase. The links in the kill chain can therefore be thought of as modular elements, with not every element present in every case.

Terminology

TTPs. We use the industry’s traditional framing of TTPs, where tactics are the highest level of observed behavior. Each tactic is broken down into a number of more specific techniques, and each technique is broken down into the most granular level of procedures.

We consider each tactic to be a separate link in the kill chain: disrupt one tactic, and you can disrupt an entire operation.

Assets. Anything that the operation controls or gains access to can be an asset. This can include both online and offline resources. Online assets include various types of social media and email accounts, but also websites, cryptocurrency wallets, and malware. Offline assets include SIM cards, bank accounts, office buildings (such as the “troll farms” exposed in Albania³⁹ and Nicaragua⁴⁰), and even office furniture (such as the beanbags that characterized one Russian troll farm).⁴¹

Information. We understand “information” in the broadest sense, to include electronic data and information about the real world. For example, a list of targets’ social media accounts, a database of compromised passwords, the movements of ships and aircraft, or the office address of a business would all count as information for the purposes of our kill chain.

Engagement. Engagement is any way an operation attempts to interact with people who are not part of it. It does not presuppose that the attempt is successful: a network like the Chinese Spamouflage network⁴² may sometimes use common hashtags to attract attention (tactic: targeted engagement; technique: posting to reach a specific audience; procedure: posting hashtags), but its posts typically received no engagement from assets outside the operation itself.

Harm. We consider “harm” to be any behavior that actually or potentially puts people at risk of physical harm, deceives or defrauds them, compromises their personal information, silences their voice, or promotes criminal activity.

We developed the Online Operations Kill Chain based on analysis of the behaviors that Meta’s threat intelligence teams regularly tackle, such as cyber espionage, influence operations, human exploitation, terrorism and organized crime, scams, and coordinated reporting and harassment. Other platforms and entities may see a scope for additional harms.

Online operation. As noted above, we use the term “online operation” as shorthand to describe a coordinated set of activities conducted by a threat actor with the apparent intent of causing harm. The kill chain is designed to analyze online operations, identify their weak points, and enable investigators to disrupt them.

The Online Operations Kill Chain

The kill chain consists of ten links. Each link represents a top-level tactic—a broad approach that threat actors use. Each tactic is broken down into more detailed techniques, which break down into yet more detailed procedures (see table 1). Procedures can be coupled with nonbehavioral metadata (such as country of origin) to produce a fine-grained picture of the operation.

Table 1. Examples of TTPs Within the Online Operations Kill Chain

Early stage of kill chain	
Tactic	Acquiring assets
Technique	Acquiring email address
Procedure	Acquiring encrypted email address from specific provider
Late stage of kill chain	
Tactic	Enabling longevity
Technique	Changing administrators
Procedure	Giving unwitting individuals administrative rights on social media assets

At ten links, the Online Operations Kill Chain is longer than most other kill chains. This is primarily because most kill chains begin with the “reconnaissance” phase. It is our position that for an operation to conduct reconnaissance, especially on social media, it most likely will have gone through other steps first (such as acquiring IP addresses, emails and/or phone numbers, and social media accounts, as well as likely disguising those assets to make them harder to detect). These “upstream” stages are reflected in our kill chain—although not all platforms or entities will be able to observe them.

Authors’ note: all case-specific examples referenced in the following sections are drawn from public reporting.

Phase 1: Acquiring Assets

This refers to any instance in which an operation acquires or sets up an asset or capability.⁴³ Such assets can range from IP and email addresses to social-media accounts and malware to physical locations in a city.

For example, as Meta reported in April 2022,⁴⁴ the hybrid cyber espionage and influence operation from Azerbaijan acquired commodity surveillanceware for Android and publicly available hash-cracking tools. An Iranian cyber-espionage operation that Meta disrupted early in 2022 created a hitherto unknown strain of malware dubbed “HilalRAT.”⁴⁵

The original troll farm, the Russian Internet Research Agency, started out in 2013 by renting office space in Saint Petersburg, and it “purchased credit card and bank account numbers from online sellers”⁴⁶ while a successor operation in early 2020 acquired a building in Ghana as a base of operations.⁴⁷ An Iranian influence operation first reported by FireEye created a number of purported news websites to spread its message.⁴⁸ Many scams register front businesses to gather and launder their proceeds.

Examples of asset acquisition within the Online Operations Kill Chain:

- Acquiring encrypted email addresses
- Acquiring social media assets
- Registering businesses
- Renting office space
- Registering web domains

Phase 2: Disguising Assets

This tactic covers any action an operation uses to make its assets look authentic. This can range from stealing profile pictures from celebrities to creating deeply backstopped personas across multiple social media platforms and websites.

For example, many operations have sought to disguise their fake accounts by giving them profile pictures likely created from freely available websites using GANs.⁴⁹ Some sexual predators pose as adolescents in their online engagements with potential victims.⁵⁰ An Iranian cyber espionage operation that Meta disrupted in July 2021⁵¹ ran cross-platform, backstopped accounts that posed as recruiters, defense and aerospace employees, journalists, medical staff, and even an aerobics instructor.⁵² Many scammers have impersonated military officers.⁵³

Asset disguise is an essentially static tactic: the threat actor selects a persona of greater or lesser sophistication and maintains it with more or less regularity. This is distinct from efforts to evade detection, described below, which are an ongoing, often repetitive practice.

Examples of asset disguise within the Online Operations Kill Chain:

- Using StyleGAN 2 profile pictures
- Impersonating real people or organizations
- Posing as fictional media outlets
- Using remote infrastructure appropriate to the target country
- Backstopping personas across multiple platforms

Phase 3: Gathering Information

This covers any effort an operation makes to gather information, whether manually or by automation. It includes not only manual or scaled cyber reconnaissance techniques, scraping, and accessing databases of stolen passwords but also using open-source registers of marine or air traffic, searching corporate registries, and viewing potential targets' social media profiles.

Much of this activity happens out of the public eye and is primarily visible to platforms, data system managers, companies, and law enforcement. For example, an agent for Chinese intelligence in the United States used “various social media sites” to research potential recruits from 2015 to 2020, according to the U.S. Department of Justice (DOJ).⁵⁴ Also according to the DOJ, the Internet Research Agency tracked “certain metrics” of American social media groups, including “the group’s size, the frequency of content placed by the group, and the level of audience engagement with that content, such as the average number of comments or responses to a post.”⁵⁵ In 2021, Meta disrupted seven providers of abusive commercial services that targeted journalists, dissidents, critics of authoritarian regimes, families of opposition, and human rights activists with surveillance-for-hire techniques.⁵⁶

Examples of information gathering within the Online Operations Kill Chain:

- Using commercially available surveillance-for-hire tools
- Using open-source flight tracking data
- Searching for targets on social media platforms
- Scraping public information
- Monitoring trending topics

Phase 4: Coordinating and Planning

This covers any method an operation uses to coordinate and plan its activity. This can include both overt and covert coordination and both manual techniques and automation.

For example, an anti-vaccine network that Meta disrupted in France and Italy in late 2021 used Telegram channels to coordinate and train people in online harassment.⁵⁷ Some of this coordination was exposed by open-source researchers.⁵⁸ Right-wing activists in 2016 were reported to be using direct message chat rooms on Twitter to coordinate their targets and use of bots.⁵⁹ Left-wing activists at the Alabama special election in 2017 used publicly viewable spreadsheets to coordinate their supporters' posting;⁶⁰ a Mexican operator showed on video how he used a spreadsheet to coordinate automated Twitter activity in 2018.⁶¹

Examples of coordination and planning within the Online Operations Kill Chain:

- Coordinating via public posts
- Training recruits in private groups
- Coordinating using encrypted apps
- Publishing lists of targets and hashtags
- Automating posting across multiple accounts

Phase 5: Testing Platform Defenses

Some operations test the limits of online detection and enforcement by sending a range of content with varying degrees of violation and observing which ones are detected.

For example, the Russian military intelligence unit that targeted Hillary Clinton's presidential campaign servers in 2016 sent test spearphishing emails as part of its preparation.⁶² Hacking groups may upload their own malware to an antivirus data website like VirusTotal to see if it would be detected. Operations that exchange or post violating content, such as hate speech or sexually explicit imagery, may post variations of the same message to see which ones are detected automatically.

Examples of defense testing within the Online Operations Kill Chain:

- Sending phishing links to operation-controlled email accounts
- Posting A/B variations of violating images
- Posting A/B variations of violating texts
- Testing own malware using publicly available tools
- Posting spam at different rates from different accounts

Phase 6: Evading Detection

Any repetitive method an operation uses to sidestep online defenses qualifies as evading detection. This can include the use of camouflaged or edited text or images and also technical measures such as routinely changing IP addresses.

For example, one method used by the anti-vaccine operation referenced above was to write the French word “vaccin” as “vaxcin” or “vaxxin” to defeat keyword detection. Journalists have reported that the Boogaloo movement sometimes used the variant spelling Boogalo to evade detection on TikTok.⁶³ A Russian operation nicknamed “Doppelganger” that spoofed the websites of European media outlets geo-restricted the fake sites so that only people in the target countries could view them.⁶⁴

Examples of evasion within the Online Operations Kill Chain:

- Using typos to obfuscate key phrases
- Geo-limiting website audiences
- Editing images
- Routing traffic through virtual private networks (VPNs) or anonymous web browsers like Tor
- Using coded language or references

Phase 7: Indiscriminate Engagement

This tactic includes any form of posting or engagement in which the operation makes no apparent effort to reach a particular audience. For example, spammers who use fake accounts to share posts to their own timelines, or operations that post content on their own websites and do not otherwise promote it, would count as indiscriminate engagers. Often, indiscriminate engagement is characterized by what operations do *not* do: an absence of any discernible efforts to reach an audience. In effect, it is a “post and pray” strategy, dropping their content onto the internet and leaving it to users to find it.

For example, the Chinese operation Spamouflage primarily posted on YouTube, Twitter, and Facebook.⁶⁵ It used large numbers of accounts to post pro-China or anti-Western videos interspersed with innocuous landscapes and sayings, but accounts often did so without any attempt—such as hashtags or @-mentions—to attract an audience. Much of the Russian operation Secondary Infektion used a “post and pray” approach—for example, posting a blog about politics in Europe on a forum dedicated to the civil service in Pakistan.⁶⁶ The Doppelganger operation sometimes made comments about the Ukraine war in response to posts about sport or fashion.

Spam operations often fall into this category too. Networks that use one fake account to post content on a social media platform and then use other fakes to share the original post to their own timelines may make the original post appear more popular than it really was, but they are not taking any meaningful steps to reach an authentic audience.

Examples of indiscriminate engagement within the Online Operations Kill Chain:

- Publishing content on web forums inappropriate to the subject matter
- Replying to posts with no relevance to the subject matter
- Publishing on operation-controlled websites only
- Posting to operation-controlled social media timelines only
- Using operation-controlled assets to comment on posts by other operation-controlled assets, where none of the assets has authentic followers

Phase 8: Targeted Engagement

Targeted engagement, by contrast, covers any sort of method an operation uses to plant its content in front of a specific audience. It can include, for example, advertising, mentioning or replying to a target account, spearphishing, or even emailing real people and trying to trick them into becoming part of the operation.

There are many examples of targeted engagement. The Russian Internet Research Agency made heavy use of ads in 2015 and 2016;⁶⁷ in 2020, it hired real people to write for it⁶⁸ and even to run ads in the United States on its behalf.⁶⁹ Russian military intelligence used social media messaging and email to communicate with people in the United States, including reporters, in 2016.⁷⁰ An Iranian operation that focused on Scottish independence in late 2021 used independence-themed hashtags on many of its posts.⁷¹ A previously unreported Iranian hacking group that Meta disrupted in early 2022 used fake “job recruitment” personas to message and email its targets.⁷² This group created fake interview and chess apps, which would only deliver the malware payload after the targets interacted with the attacker in real time. In 2021, Google revealed North Korean actors posed as security researchers to lure other researchers into sharing vulnerabilities and exploit code.⁷³

Targeted engagement is an important late-stage TTP for researchers, because it is the area where operations likely show the most unique combination of approaches. For journalists and researchers, this is also essential security awareness training to recognize when they or their colleagues become the targets.

Examples of targeted engagement within the Online Operations Kill Chain:

- Running ads
- Using hashtags appropriate to the target audience
- Emailing potential victims or recruits
- Submitting operation material to authentic news outlets
- Directing harassment groups to specific people or posts

Phase 9: Compromising Assets

An operation that attempts to access or take over accounts or information is considered to be compromising assets. Espionage actors are the primary culprits here, but scammers and influence operations can also compromise assets under some circumstances.

Social media asset compromise can cover, for example, password spraying, spearphishing, a variety of social engineering techniques, device compromise, and access via email compromise, as in the case of the espionage and influence operation known as Ghostwriter.⁷⁴ It can also cover incidents when threat actors convince the administrators of pages or groups to make them administrators, too, and then use their new privileges to remove the other administrators from the page or group in question.⁷⁵ And it can cover compromises of third-party apps, which give the threat actor access to high-profile accounts.⁷⁶

Examples of compromise within the Online Operations Kill Chain:

- Phishing email login credentials
- Using compromised email accounts to access social media accounts
- Socially engineering victims to hand over credentials
- Acquiring administrative privileges on social media assets
- Installing malware on victim servers

Phase 10: Enabling Longevity

Finally, operations that take steps to survive takedown, or to prolong their activity after exposure, are considered to be enabling longevity. Many publicly documented operations have responded to disruption by attempting to adapt their TTPs and restore their presence on different platforms: this is why one use of the kill chain can be to compare different stages of the same operation, to analyze any forced adaptation measures and develop countermeasures.

For example, Spamouflage responded to the takedown of one of its Twitter personas (named 李若水francisw) by acquiring preexisting accounts on the platform, giving them the persona's name and profile picture, and returning to posting with the explicit message, "This is my new account."⁷⁷ When Meta blocked the first set of spoofed domains created by Doppelganger, the operation created hundreds of new domains to try to redirect people to the spoofed sites.⁷⁸ An Iranian operation known as IUVM responded to the loss of its social media assets by creating new fakes to spread its imagery.⁷⁹ After Russian military intelligence's "Alice Donovan" persona was exposed, it emailed at least one outlet that had published its work to falsely claim that "she" had deleted "her" Facebook account, but the account continued posting on Twitter.⁸⁰ As the latter example shows, operations may also spread themselves across platforms, partly in the hope that at least some accounts may evade enforcement.

Attempts to prolong the longevity of an operation can take unusual forms. In 2018, the Internet Research Agency had approximately one hundred Instagram accounts taken down shortly before the U.S. midterm elections. It responded by falsely claiming that those accounts were only the tip of the iceberg, and its operation had already thrown the elections, engaging in what we call "perception hacking." The attempt was met with ridicule, but it remains an example of trying to turn a takedown into a communications opportunity.⁸¹

Examples of enabling longevity within the Online Operations Kill Chain:

- Replacing disabled accounts with new ones using the same persona
- Changing email addresses
- Creating new web domains that redirect to old ones
- Deleting logs and other evidence
- Weaponizing a disruption to claim that it was part of the plan all along

After longevity: The daisy-chain effect. One recurring question when investigating the efforts of particularly persistent threat actors is: at what point should sufficiently determined persistence be considered a new operation? Many of the more persistent threat actors represent what could be thought of as a daisy-chain effect, in which the late-stage elements of one operation segue into the early-stage elements of a new one, and any distinction between the two is largely arbitrary.

For the sake of practicality, we consider that an operation can be treated as “new” if it changes the majority of its procedures in the first phases of the kill chain: asset acquisition and disguise. For example, a harassment network that reconstitutes after disruption by setting up accounts on the same IP addresses and reusing the visual branding of its first iteration would not count as a new operation. By contrast, when individuals associated with past activity by the Internet Research Agency began operating in Ghana in early 2020, they used entirely new physical and online infrastructure, disguised their operation as a local nongovernmental organization, and created a website and blogs—as well as social media accounts—to backstop the deception.⁸² This showed enough variation to qualify as a new operation.

Appendix: Case Studies—The Online Operations Kill Chain in Use

To illustrate how the kill chain can be used, the following case studies apply the Online Operations Kill Chain to operations that have been publicly reported with unusual detail: the hacking and leaking of Clinton campaign emails in 2016 by Russian military intelligence (the Main Directorate of the General Staff of the Armed Forces, or GRU) known as “DCLeaks”; the “PeaceData” website run by the Internet Research Agency in 2020; and the “V_V” anti-vaccine harassment movement that Meta took down in 2021.

The main sources for the GRU’s hack-and-leak operation in 2016 are the U.S. DOJ’s indictment of the suspected hackers⁸³ and its redacted report into Russian interference.⁸⁴ CounterPunch’s investigation into the “Alice Donovan” persona is a trove of information around “her” publishing activity.⁸⁵ Sources for the PeaceData operation include the original takedown announcements by Facebook⁸⁶ and Twitter;⁸⁷ the simultaneous report by Graphika based on the takedown;⁸⁸ and victim testimonies published by Reuters,⁸⁹ the Daily Beast,⁹⁰ the New York Times,⁹¹ and New Zealand news site newsroom.co.nz.⁹² The main sources for the V_V takedown are Meta’s takedown announcement⁹³ and the in-depth research conducted by Graphika.⁹⁴

The level of detail in the DOJ’s reporting gives us a rare opportunity to include in a public analysis details that would typically be private or inaccessible, such as server acquisition, financial transactions, and recruitment emails. The PeaceData and V_V cases give a more typical illustration of what can be achieved with open-source methods.

DCLeaks and Alice Donovan

Acquiring assets

- Setting up email addresses (yandex.com, mail.com, gmail.com, aol.fr)
- Leasing server in target country
- Leasing server in third country (Malaysia)
- Leasing computer in target country
- Acquiring cryptocurrency wallet
- Acquiring VPN account
- Acquiring cloud computer account
- Acquiring link-shortening account
- Acquiring social media accounts (Facebook, Twitter, Pinterest)
- Creating malware (X-Agent, X-Tunnel)
- Registering websites
- Setting up blog
- Setting up remote middleman server

Disguising assets

- Stealing profile pictures
- Creating fictional personas (Alice Donovan, DCLeaks, Guccifer 2.0)
- Backstopping personas across platforms (Facebook, Twitter, Pinterest, websites, blogs)
- Attributing own activity to external organization (claiming DCLeaks was a “Wikileaks sub-project”)
- Spoofing sender email address in spearphishing attacks
- Creating phishing domains resembling real ones (accounts-qooqle.com, account-gooogle.com)
- Creating email address one letter away from real person’s name

Gathering information

- Researching victims on social media
- Searching for open-source information about victims’ computer networks
- Querying victim IP configurations to identify connected devices

- Searching victim devices for keywords in files
- Searching for translations
- Copying articles by real authors

Coordinating and planning

- Coordinating through military chain of command
- Coordinating between distinct units (cyber units 26165 and 74455)

Testing defenses

- Testing malware ability to connect to target
- Testing ability to compress and exfiltrate data from target

Evading detection

- Using link-shortening tools to obfuscate malware links
- Using middleman server to obfuscate data exfiltration
- Using compression tools to conceal scale of data exfiltration
- Registering web domain under privacy protection

Indiscriminate engagement

- Posting content on a blog hosted by WordPress

Targeted engagement

- Sending malware to spearphishing targets by email
- Submitting articles to news websites by email
- Sending hacked content to unwitting individuals by email
- Contacting news websites by direct message
- Sending hacked content to unwitting individuals by direct message
- Posting hacked content on password-protected site
- Publishing hacked content on website on a daily basis
- Promoting hacked content on social media
- Laundering hacked content through external organization
- Curating and copying content written by genuine authors

Compromising assets

- Spearphishing target credentials
- Using stolen credentials to access victim server
- Installing malware on victim server
- Logging keystrokes
- Taking screenshots
- Exfiltrating data via middleman server

Enabling longevity

- Deleting logs and files
- Searching for open-source releases about the hackers' tools
- Replacing phishing infrastructure with new phishing site (actblues[.]com)
- Using fake persona to deny public attribution (Guccifer 2.0)
- Engaging with editors after exposure to proclaim innocence (Alice Donovan)
- Claiming to have self-deleted social media accounts that were actually taken down, arguing this was “for safety reasons” (Alice Donovan)
- Removing bylines of exposed fake personas from websites controlled by the operation, but leaving the articles up (Alice Donovan/Inside Syria Media Centre)

PeaceData

Acquiring assets

- Setting up email addresses on encrypted domain (Proton Mail)
- Setting up email addresses on own domain (peacedata.net)
- Acquiring online payment account (PayPal)
- Acquiring social media accounts (Facebook, Twitter, WhatsApp, LinkedIn, UpWork, Guru)
- Acquiring inauthentic friends/followers
- Setting up websites (peacemonitor.com, peacedata.net)

Disguising assets

- Using GAN-generated profile pictures
- Running inauthentic media brand
- Running fake personas
- Pretending to be located in third countries
- Backstopping personas across platforms (Facebook, Twitter, LinkedIn, website, author bylines, emails)
- Giving fake personas specific roles within fake brand (such as recruiting, editor, or deputy editor)

Gathering information

- Copying news articles from authentic sites
- Searching for freelance contributors on social media
- Searching for job-listing sites appropriate to target audience

Coordinating and planning

- Coordinating using encrypted email (Proton Mail)
- Coordinating using encrypted messaging (WhatsApp)
- Creating fake publishing partnership with external websites

Evading detection

- Recruiting unwitting contributor in America to run political Facebook ads
- Recruiting unwitting native-language authors
- Recruiting professional translator
- Moving communications from social media messaging to email

Indiscriminate engagement

- No evidence (engagement was primarily targeted)

Targeted engagement

- Running ads for freelance writers on job forums
- Cold messaging potential contributors on LinkedIn
- Emailing potential contributors
- Direct messaging potential contributors on social media
- Recruiting contributors in target countries
- Paying contributors via PayPal
- Sharing links into politically aligned Facebook groups
- Asking unwitting contributors to amplify publications to their own networks
- Adding political slant to some articles

Compromising assets

- No evidence

Enabling longevity

- Giving unwitting individuals admin rights on social media assets
- Denying exposure in public statement
- Denying exposure in private communications to contributors

V_V

Acquiring assets

- Acquiring emails
- Acquiring phone numbers
- Acquiring authentic social media accounts (Facebook, Telegram, Instagram, YouTube, TikTok, VKontakte)
- Acquiring duplicate social media accounts
- Acquiring inauthentic social media accounts

Disguising assets

- Branding assets with V_V logo

Coordinating and planning

- Creating public hierarchy within organization
- Training new recruits using social media posts
- Coordinating harassment in private channels
- Coordinating on encrypted messaging apps (WhatsApp, Signal)
- Coordinating posting assignments (for example, memes, links and videos)
- Coordinating via shared hashtags
- Allocating a rank/number to each member

Evading detection

- Scrambling letters in key words (“vaccine”/“vaccine”)
- Replacing letters with numbers in key words (“v4ccine”)
- Replacing letters with emojis in key words (√ instead of V)
- Switching channels from public to private and back at set times

Indiscriminate engagement

- Distributing printed flyers through residents’ physical mailboxes

Targeted engagement

- Mass down-voting of targets’ posts
- Mass commenting on targets’ posts
- Mass posting hashtags
- Defacing targets’ personal photos with Nazi imagery
- Inviting users of other platforms to join Telegram
- Tagging friends to attract them to branded content
- Mass voting on online polls
- Graffiti on target buildings

Compromising assets

- Mass booking genuine vaccination appointments and then cancelling them at the last minute

Enabling longevity

- Operating across platforms to take advantage of different enforcement regimes

About the Authors

Ben Nimmo is Meta's global threat intelligence lead. He was a co-founder of the Atlantic Council's Digital Forensic Research Lab (DFRLab), and later served as Graphika's first head of investigations. He has helped to expose foreign election interference in the United States, United Kingdom and France; documented troll operations in Asia, Africa, Europe and the Americas; and been declared dead by an army of Twitter bots. A graduate of Cambridge University, he speaks French, German, Russian, and Latvian, among other languages.

Eric Hutchins is a security engineer investigator on Meta's influence operations team. During his nineteen-year career at Lockheed Martin, he co-authored the seminal Intrusions Kill Chain white paper, forged partnerships across industry and government, and founded the premiere enterprise network defense and insider threat team, LM-CIRT. He was the youngest engineer ever to achieve the seniormost rank of Lockheed Martin fellow. In 2021, he brought his net defense mindset to Meta and to a different problem space of countering influence operations.

Notes

- 1 For example, the U.S. State Department's Global Engagement Center, described online at <https://www.state.gov/bureaus-offices/under-secretary-for-public-diplomacy-and-public-affairs/global-engagement-center>, and the French government's VIGINUM, described online at http://www.sgdsn.gouv.fr/communiqués_presse/viginum-le-dispositif-de-protection-contre-les-ingerences-numeriques-etrangeres-se-structure/.
- 2 For example, the Atlantic Council's Digital Forensic Research Lab, online at <https://medium.com/dfirlab>, and the Australian Security Policy Institute, online at www.aspi.org.au.
- 3 Notably Graphika, <https://graphika.com>, and FireEye, www.fireeye.com.
- 4 For example, the Oxford Internet Institute, online at <https://www.oii.ox.ac.uk>; the Stanford Internet Observatory, online at <https://cyber.fsi.stanford.edu/io>; EU DisinfoLab, <https://www.disinfo.eu/>; and Cardiff University's Open Source Communications, Analytics and Research Development Centre, <https://www.cardiff.ac.uk/crime-security-research-institute/research/projects/open-source-communications-analytics-and-research-development-centre>.
- 5 See, for example, Meta's Coordinated Inauthentic Behavior archives, online at <https://about.fb.com/news/tag/coordinated-inauthentic-behavior/>; Twitter's archive of information operations, online at <https://transparency.twitter.com/en/reports/information-operations.html>; and Google's Threat Analysis Group blogs, online at <https://blog.google/threat-analysis-group/>.
- 6 "Online Safety Bill," UK House of Commons, 2022, <https://bills.parliament.uk/bills/3137/publications>.
- 7 "2022 Strengthened Code of Practice on Disinformation," European Commission, June 16, 2022, <https://digital-strategy.ec.europa.eu/en/library/2022-strengthened-code-practice-disinformation>. The strengthened code was published while this paper was being written and calls for the development of a "cross-service understanding of manipulative behaviours, actors and practices" not permitted on social media. The Online Operations Kill Chain can be used to answer that call.
- 8 "The Digital Services Act: Ensuring a Safe and Accountable Online Environment," European Commission, December 15, 2020, https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/digital-services-act-ensuring-safe-and-accountable-online-environment_en.

- 9 Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," Lockheed Martin, 2011, <https://www.lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>.
- 10 Paul Pols, "The Unified Kill Chain," Fox-IT, 2017, <https://www.unifiedkillchain.com/assets/The-Unified-Kill-Chain.pdf>.
- 11 "MITRE ATT&CK," MITRE Corporation, 2015, <https://attack.mitre.org/>.
- 12 Sergio Caltagirone, Andrew Pendergast, and Christopher Betz, "Diamond Model of Intrusion Analysis," Technical Report ADA586960, Center for Cyber Threat Intelligence and Threat Research, July 5, 2013, <https://www.threatintel.academy/diamond/>.
- 13 David J. Bianco, "The Pyramid of Pain," *Enterprise Detection & Response* (blog), March 1, 2013, <http://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>.
- 14 Bruce Schneier, "Influence Operations Kill Chain," *Schneier on Security* (blog), August 2019, https://www.schneier.com/blog/archives/2019/08/influence_opera.html.
- 15 Clint Watts, "Advanced Persistent Manipulators, Part Three: Social Media Kill Chain," Alliance for Securing Democracy, July 22, 2019, <https://securingdemocracy.gmfus.org/advanced-persistent-manipulators-part-three-social-media-kill-chain/>.
- 16 Katerina Sedova, Christine McNeill, Aurora Johnson, Aditi Joshi, and Ido Wulkan, "AI and the Future of Disinformation Campaigns, Part 1: The RICHDATA Framework," Georgetown Center for Security and Emerging Technology, December 2021, <https://cset.georgetown.edu/publication/ai-and-the-future-of-disinformation-campaigns/>.
- 17 John F. Gray and Sara-Jayne Terp, "Misinformation: We're Four Steps Behind Its Creators," Credibility Coalition Misinfosec Working Group, November 2019, <https://cyber.harvard.edu/sites/default/files/2019-11/Comparative%20Approaches%20to%20Disinformation%20-%20John%20Gray%20Abstract.pdf>. The AMITT framework was updated and named the DISARM framework in 2022: <https://github.com/DISARMFoundation/DISARMframeworks/>.
- 18 Photon Research Team, "The Account Takeover Kill Chain: A Five Step Analysis," *Cybercrime and Dark Web Research* (blog), Digital Shadows, July 30, 2019, <https://www.digitalsadows.com/blog-and-research/the-account-takeover-kill-chain-a-five-step-analysis/>.
- 19 John Johnson, "Cyber Fraud Kill Chain," TechStrong TV interview, video, 11:45, September 23, 2021, <https://digitalanarchist.com/videos/interviews/cyber-fraud-kill-chain-optim>.
- 20 We use the term "online operation" as shorthand to describe a coordinated set of activities conducted by a threat actor with the apparent intent of causing harm. Much has already been written about how best to define harmful behaviors; that discussion is beyond the scope of this paper.
- 21 For example, "2022 Strengthened Code of Practice on Disinformation," European Commission.
- 22 For example, Schneier, "Influence Operations Kill Chain."
- 23 For example, Gray and Terp, "Misinformation: We're Four Steps Behind Its Creators."
- 24 See, for example, "How to Identify Misinformation, Disinformation, and Malinformation," Canadian Centre for Cyber Security, February 2022, <https://cyber.gc.ca/en/guidance/how-identify-misinformation-disinformation-and-malinformation-itsap00300>.
- 25 For example, Nathaniel Gleicher, Margarita Franklin, David Agranovich, Ben Nimmo, Olga Belogolova, and Mike Torrey, "Threat Report: The State of Influence Operations 2017-2020," Facebook, May 2021, <https://about.fb.com/wp-content/uploads/2021/05/IO-Threat-Report-May-20-2021.pdf>.
- 26 Hutchins, Cloppert, and Amin, "Intelligence-Driven Computer Network Defense."
- 27 Gray and Terp, "Misinformation: We're Four Steps Behind Its Creators."

- 28 Lee Foster, Sam Riddell, David Mainor, and Gabby Roncone, “Ghostwriter’ Influence Campaign: Unknown Actors Leverage Website Compromises and Fabricated Content to Push Narratives Aligned With Russian Security Interests,” Mandiant, July 28, 2020, <https://www.mandiant.com/resources/ghostwriter-influence-campaign/>.
- 29 Ben Nimmo, David Agranovich, and Nathaniel Gleicher, “Adversarial Threat Report, Q1 2022,” Meta, April 7, 2022, https://about.fb.com/wp-content/uploads/2022/04/Meta-Quarterly-Adversarial-Threat-Report_Q1-2022.pdf.
- 30 “Indictment in the case of United States of America v. Viktor Borisovich Netyksho et al.,” U.S. Department of Justice Case 1:18-cr-00215-ABJ, July 13, 2018, <https://www.justice.gov/file/1080281/download/>.
- 31 We have optimized this kill chain framework to analyze operations that aim at triggering a degree of interaction or engagement among human targets. This is designed to complement the work done to develop various kill chains that focus on machine-targeted operations in the cybersecurity space.
- 32 One of the most striking changes in online operations of the past five years has been the adoption by threat actors of profile pictures created using Generative Adversarial Networks (GAN) apparently sourced from public websites. The first observation of the use of this technology at scale is at Ben Nimmo, C. Shawn Eib, L. Tamora, Kate Johnson, Ian Smith, Eto Buziashvili, Alyssa Kann, Kanishk Karan, Esteban Ponce de León Rosas, and Max Rizzuto, “Operation FFS: Fake Face Swarm,” Graphika, December 20, 2019, <https://graphika.com/reports/operationffs-fake-face-swarm/>. By April 2022, all four networks that Meta took down for coordinated inauthentic behavior used the same procedure (see Nimmo, Agranovich, and Gleicher, “Adversarial Threat Report, Q1 2022”).
- 33 The kill chain relies on post hoc analysis of online operations to describe the detailed sequence of actions that they undertake. As such, its immediate utility is as a framework to analyze, compare, and share information about existing activity. However, insights into existing operations can also be used to inform defenses against future operations—for example, by identifying hitherto unknown TTPs and prioritizing the development of ways to detect them.
- 34 For example, the Russian operation Secondary Infektion experimented with fictitious personas in 2014–2016, before shifting to a focus on single-use burner accounts. For a description of the early personas, see Ben Nimmo, Camille François, C. Shawn Eib, Lea Ronzaud, Rodrigo Ferreira, Chris Hernon, and Tim Kostelancik, “Secondary Infektion: Early Experiments With Personas,” Graphika, June 16, 2020, <https://secondaryinfektion.org/report/early-experiments-with-personas/>.
- 35 Different recipients will have different needs when it comes to the granularity of what is shared: for example, industry threat disruption teams are likely to be most interested in detailed selectors, whereas regulators are more likely to be interested in persistent and durable behaviors.
- 36 Nathaniel Gleicher, Ben Nimmo, David Agranovich and Mike Dvilyanski, “Adversarial Threat Report,” Meta, December 1, 2021, <https://about.fb.com/wp-content/uploads/2021/12/Metas-Adversarial-Threat-Report.pdf>.
- 37 For example, Graphika’s work on Russian operation Secondary Infektion in Nimmo, François, Eib, Ronzaud, Ferreira, Hernon, and Kostelancik, “Secondary Infektion,” and Lockheed Martin’s exposure of a Chinese zero-day exploit reportedly based on a National Security Agency original, in Andy Greenberg, “China Hijacked an NSA Hacking Tool in 2014—and Used It for Years,” *Wired*, February 22, 2021, <https://www.wired.com/story/china-nsa-hacking-tool-epme-hijack/>.
- 38 For example, Clarissa Ward, Katie Polglase, Sebastian Shukla, Gianluca Mezzofiore, and Tim Lister, “Russian Election Meddling Is Back—Via Ghana and Nigeria—and In Your Feeds,” CNN, March 12, 2020, <https://edition.cnn.com/2020/03/12/world/russia-ghana-troll-farms-2020-ward/index.html>.
- 39 Ethan Fecht, Ben Nimmo, and the Facebook Threat Intelligence team, “In-depth Analysis of the Albania-based Network,” Facebook, April 2021, <https://about.fb.com/wp-content/uploads/2021/04/March-2021-CIB-Report.pdf>.
- 40 Luis Fernando Alonso, Ben Nimmo, and the Meta Threat Intelligence team, “October 2021 Coordinated Inauthentic Behavior Report,” Meta, November 2021, <https://about.fb.com/wp-content/uploads/2021/11/October-2021-CIB-Report.pdf>.

- 41 Kseniya Klochkova, “«Вы ведь не верите, что это настоящие отзывы?» Как «Фонтанка» заглянула на передовую информационных фронтов Z [You don't really believe that they're real reviews? How 'Fontanka' looked at the front line of the Z information fronts],” Fontanka, March 21, 2022, <https://www.fontanka.ru/2022/03/21/70522490/>.
- 42 Ben Nimmo, C. Shawn Eib and L. Tamora, “Spamouflage,” Graphika, September 25, 2019, <https://graphika.com/reports/spamouflage/>.
- 43 Investigators at social media platforms and service providers may find it useful to distinguish between assets acquired from other providers and assets created on their own platform. However, such a distinction would create ambiguity for the broader open-source community around what constitutes “acquisition” as opposed to “creation.”
- 44 Nimmo, Agranovich and Gleicher, “Adversarial Threat Report, Q1 2022.”
- 45 Nimmo, Agranovich and Gleicher, “Adversarial Threat Report, Q1 2022.”
- 46 “Indictment in the case of United States of America v. Internet Research Agency et al.,” U.S. Department of Justice case 1:18-cr-00032-DLF, February 16, 2018, <https://www.justice.gov/file/1035477/download>.
- 47 Ward, Polglase, Shukla, Mezzofiore, and Lister, “Russian Election Meddling Is Back.”
- 48 “Suspected Iranian Influence Operation,” FireEye, August 2018, <https://www.mandiant.com/resources/report-suspected-iranian-influence-operation>.
- 49 Nimmo, Eib, Tamora, Johnson, Smith, Buziashvili, Kann, Karan, Ponce de León Rosas, and Rizzuto, “#OperationFFS: Fake Face Swarm.”
- 50 Janis Wolak, David Finkelhor, Kimberly J. Mitchell, and Michele L. Ybarra, “Online ‘Predators’ and Their Victims,” *American Psychologist* 63, no. 2, February–March 2008, <https://www.apa.org/pubs/journals/releases/amp-632111.pdf>.
- 51 Mike Dvilyanski and David Agranovich, “Taking Action Against Hackers in Iran,” Facebook, July 15, 2021, <https://about.fb.com/news/2021/07/taking-action-against-hackers-in-iran/>.
- 52 Joshua Miller, Michael Raggi, and Crista Giering, “I Knew You Were Trouble: TA456 Targets Defense Contractor With Alluring Social Media Persona,” Proofpoint, July 28, 2021, <https://www.proofpoint.com/us/blog/threat-insight/i-knew-you-were-trouble-ta456-targets-defense-contractor-alluring-social-media/>.
- 53 See, for example, Amelia Shaw and Scott Campbell, “Widow Who Thought She'd Found Love Online Scammed Out of Savings By Conman Pretending to Be US General,” *Daily Mirror*, August 12, 2016, <https://www.mirror.co.uk/news/uk-news/widow-who-thought-she-d-found-love-8620732>.
- 54 “Singaporean National Pleads Guilty to Acting in the United States as an Illegal Agent of Chinese Intelligence,” U.S. Department of Justice, press release, July 24, 2020, <https://www.justice.gov/opa/pr/singaporean-national-pleads-guilty-acting-united-states-illegal-agent-chinese-intelligence/>.
- 55 U.S. Department of Justice, “Indictment in the case of United States of America v. Internet Research Agency et al.”
- 56 Mike Dvilyanski, David Agranovich, and Nathaniel Gleicher, “Threat Report on the Surveillance-for-Hire Industry,” Meta, December 16, 2022, <https://about.fb.com/wp-content/uploads/2021/12/Threat-Report-on-the-Surveillance-for-Hire-Industry.pdf>.
- 57 Gleicher, Nimmo, Agranovich and Dvilyanski, “Adversarial Threat Report,” Meta, December 1, 2021, <https://about.fb.com/wp-content/uploads/2021/12/Metas-Adversarial-Threat-Report.pdf>.
- 58 The Graphika Team, “Viral Vendetta,” December 1, 2021, https://public-assets.graphika.com/reports/graphika_report_viral_vendetta.pdf.
- 59 Shawn Musgrave, “The Secret Twitter Rooms of Trump Nation,” *Politico*, August 9, 2017, <https://www.politico.eu/article/twitter-donald-trump-the-secret-twitter-rooms-of-trump-nation/>.
- 60 Ben Nimmo, “#ElectionWatch: Alabama Twitter War,” DFRLab, November 17, 2017, <https://medium.com/dfrlab/electionwatch-alabama-twitter-war-47b34ae89c50/>.

- 61 Ryan Broderick and Ínigo Arredondo, “Meet the 29-Year-Old Trying to Become the King Of Mexican Fake News,” BuzzFeed, June 28, 2018, <https://www.buzzfeednews.com/article/ryanhatesthis/meet-the-29-year-old-trying-to-become-the-king-of-mexican#.hnG3ZOL62d>.
- 62 U.S. Department of Justice, “Indictment in the case of United States of America v. Viktor Borisovich Netyksho et al.”
- 63 Alex Kaplan, “TikTok Is Full of ‘Boogaloo’ Videos Even Though It Prohibits Content From ‘Dangerous Individuals and Organizations,’” Media Matters, June 5, 2020, <https://www.mediamatters.org/tiktok/tiktok-full-boogaloo-videos-even-though-it-prohibits-content-dangerous-individuals-and/>.
- 64 Alexandre Alaphilippe, Gary Machado, Raquel Miguel, and Francesco Poldi, “Doppelganger: Media Clones Serving Russian Propaganda,” EU DisinfoLab, September 27, 2022, <https://www.disinfo.eu/doppelganger>.
- 65 Nimmo, Eib, and Tamora, “Spamouflage.”
- 66 Nimmo, François, Eib, Ronzaud, Ferreira, Hernon, and Kostelancik, “Exposing Secondary Infektion,” Graphika, June 16, 2020, <https://secondaryinfektion.org/>.
- 67 U.S. Department of Justice, “Indictment in the case of United States of America v. Internet Research Agency et al.”
- 68 Jack Stubbs, “Duped by Russia, Freelancers Ensnared in Disinformation Campaign by Promise of Easy Money,” Reuters, September 2, 2020, <https://www.reuters.com/article/us-usa-election-facebook-russia-idUSKBN25T35E>.
- 69 Adam Rawnsley, “She Was Tricked by Russian Trolls—and It Derailed Her Life,” Daily Beast, September 6, 2020, <https://www.thedailybeast.com/she-was-tricked-by-russian-trollsand-it-derailed-her-life/>.
- 70 Special Counsel Robert S. Mueller III, “Report on the Investigation Into Russian Interference in the 2016 Presidential Election,” United States Department of Justice, March 2019, <https://www.justice.gov/archives/sco/file/1373816/download>.
- 71 “December 2021 Coordinated Inauthentic Behavior Report,” Meta, January 20, 2022, <https://about.fb.com/wp-content/uploads/2022/01/December-2021-Coordinated-Inauthentic-Behavior-Report-2.pdf>.
- 72 Nimmo, Agranovich, and Gleicher, “Adversarial Threat Report, Q1 2022.”
- 73 Adam Weidemann, “New Campaign Targeting Security Researchers,” Google Threat Analysis Group, January 25, 2021, <https://blog.google/threat-analysis-group/new-campaign-targeting-security-researchers>.
- 74 Gabriella Roncone, Alden Wahlstrom, Alice Revelli, David Mainor, Sam Riddell, and Ben Read, “UNC1151 Assessed With High Confidence to Have Links to Belarus, Ghostwriter Campaign Aligned With Belarusian Government Interests,” Mandiant, November 16, 2021, <https://www.mandiant.com/resources/unc1151-linked-to-belarus-government>.
- 75 Craig Timberg, “The Facebook Page ‘Vets for Trump’ Was Hijacked by a North Macedonian Businessman. It Took Months for the Owners to Get It Back,” *Washington Post*, September 17, 2019, <https://www.washingtonpost.com/technology/2019/09/17/popular-facebook-page-vets-trump-seemed-be-place-former-military-months-macedonians-controlled-it/>.
- 76 Jon Russell, “Prominent Twitter Accounts Compromised After Third-party App Twitter Counter Hacked,” TechCrunch, March 15, 2017, <https://techcrunch.com/2017/03/15/twitter-counter-hacked/>.
- 77 Ben Nimmo, Ira Hubert, and Yang Cheng, “Spamouflage Breakout,” Graphika, February 2021, https://public-assets.graphika.com/reports/graphika_report_spamouflage_breakout.pdf.
- 78 As of October 6, 2022, Meta had blocked and exposed over 250 of these redirect domains. See the updated report and list of domains at Ben Nimmo and Mike Torrey, “Taking Down Coordinated Inauthentic Behavior From Russia and China,” Meta, September 2022, https://about.fb.com/wp-content/uploads/2022/09/CIB-Report_-China-Russia_Sept-2022-1-1.pdf.
- 79 Ben Nimmo, Camille François, C. Shawn Eib, and Léa Ronzaud, “Iran’s IUVM Turns to Coronavirus,” Graphika, April 2020, https://public-assets.graphika.com/reports/Graphika_Report_IUVM_Turns_to_Coronavirus.pdf.

- 80 Jeffrey St. Clair and Joshua Frank, "Go Ask Alice: the Curious Case of 'Alice Donovan,'" Counterpunch, December 25, 2017, <https://www.counterpunch.org/2017/12/25/go-ask-alice-the-curious-case-of-alice-donovan-2/>.
- 81 Casey Michel, "Russian Trolling Ahead of the Midterm Elections Is a Mixture of the Weird and the Pathetic," ThinkProgress, November 6, 2018, <https://archive.thinkprogress.org/usaira-russian-trolling-ahead-of-the-midterm-elections-is-a-mixture-of-the-weird-and-the-pathetic-c64b0d120ae5/>.
- 82 Ward, Polglase, Shukla, Mezzofiore, and Lister, "Russian Election Meddling Is Back."
- 83 U.S. Department of Justice, "Indictment in the case of United States of America v. Viktor Borisovich Netyksho et al."
- 84 Mueller, "Investigation Into Russian Interference in the 2016 Presidential Election."
- 85 St. Clair and Frank, "Go Ask Alice."
- 86 Facebook, "August 2020 Coordinated Inauthentic Behavior Report."
- 87 Twitter Safety (@TwitterSafety), "We suspended five Twitter accounts for platform manipulation that we can reliably attribute to Russian state actors. As standard, they will be included in updates to our database of information operations in the coming weeks to empower academic research," Twitter, 1:30 p.m., September 1, 2020, <https://twitter.com/TwitterSafety/status/1300848632120242181>.
- 88 Ben Nimmo, Camille François, C. Shawn Eib and Léa Ronzaud, "IRA Again: Unlucky Thirteen," Graphika, September 1, 2020, https://public-assets.graphika.com/reports/graphika_report_ira_again_unlucky_thirteen.pdf.
- 89 Stubbs, "Duped by Russia."
- 90 Rawnsley, "She Was Tricked by Russian Trolls."
- 91 Sheera Frenkel, "A Freelance Writer Learns He Was Working for the Russians," *New York Times*, September 2, 2020, <https://www.nytimes.com/2020/09/02/technology/peacedata-writer-russian-misinformation.html>.
- 92 Laura Walters, "I Was Part of a Russian Meddling Campaign," Newsroom.co.nz, September 4, 2020, <https://www.newsroom.co.nz/i-was-part-of-a-russian-meddling-campaign>.
- 93 Gleicher, Nimmo, Agranovich, and Dvilyanski, "Adversarial Threat Report."
- 94 The Graphika Team, "Viral Vendetta."

Carnegie Endowment for International Peace

The Carnegie Endowment for International Peace is a unique global network of policy research centers around the world. Our mission, dating back more than a century, is to advance peace through analysis and development of fresh policy ideas and direct engagement and collaboration with decisionmakers in government, business, and civil society. Working together, our centers bring the inestimable benefit of multiple national viewpoints to bilateral, regional, and global issues.

Partnership for Countering Influence Operations

Influence operations are a complex threat, and the community combating them—academics, social platforms, think tanks, governments—is broad. The goal of the Partnership for Countering Influence Operations (PCIO) is to foster evidence-based policymaking to counter threats in the information environment. Key roadblocks as found in our work include the lack of: transparency reporting to inform what data is available for research purposes; rules guiding how data can be shared with researchers and for what purposes; and an international mechanism for fostering research collaboration at-scale.



 **CARNEGIE**
ENDOWMENT FOR
INTERNATIONAL PEACE

CarnegieEndowment.org