Eye on Security

The Ghostwriter Campaign and Trends in Disinformation Today

Transcript

Luke McNamara:

Welcome to another episode of the Eye on Security podcast. I am your host, Luke McNamara. With me today to discuss the world of disinformation that this team focuses on tracking is our senior manager of information operations analysis, Lee Foster. Lee, welcome to the podcast.

Lee Foster:

Thanks for having me.

Luke McNamara:

So to kick this off – and we definitely want to get into talking about not just what your team is focused on seeing and tracking this year when it comes to the trends in the space, but certainly the Ghostwriter report that you recently released. Before we even get into that, though, I think it would be useful to hear kind of what is your framing of what you track around information operations? What's the remit and focus of your team?

Lee Foster:

Yes. So, information operations is a huge field, and different people have different definitions for what IO encompasses. There's not really a unified definition for information operations across all of the people that do research in that area. But we focus on what I call a narrow subset of information operations, which is really online influence operations. And often these are supported by what you might term traditional cyber threat activities. So things like intrusions and defacements, and so on. But obviously, this is not always the case. And we also track your more generic social media-driven influence operations online, as well.

Luke McNamara:

So when you're looking at the threat activity that we're seeing this year and how that's changed over the years, what are some of the noticeable differences that we're seeing now across some of the threat activity that we see for some of the campaigns that we've attributed, at least this sort of messaging that we see. Or potentially in some cases where we're seeing new threat actors, what are the TDPs, what are we currently seeing in terms of threat activity in this space?

Lee Foster:

So probably the biggest kind of trend, for lack of a better term, is just the sheer number of players that have gone into the space, certainly since I built this team in early 2016, tracking kind of Russian activity targeting the US, and we were obviously well in place for when the DC leaks and DNC-related activity took place. But since then we've really unmasked such a diverse array of different players in this space. This includes both nation states. We were the first, for example, to expose Iran's activity. But it also includes non-state actors, as well, given so many of kind of the IO techniques are readily deployable with kind of minimal barriers to entry.

So we see a lot of kind of non-state actors, as well, engaging in this type of activity. Beyond that there's a lot of other kind of developments in the space. Increasing incidents of outsourcing IO campaigns by state actors to effectively private entities, PR firms, and so on, that possess kind of a nuanced understanding of a particular information environment and have to kind of know how in terms of how to influence a particular audience. So we see outsourcing as kind of a growing trend in this space. And we've seen increasing kind of operational security from state-backed actors.

Some of this activity is getting more and more difficult to detect, as the operators get better at masking what they are doing. This is particularly true for social media platforms. They have experienced kind of greater difficulty in identifying or directly attributing to state actors certain campaigns because of the operational security that the actors are employing.

We've also for awhile now seen the quite extensive use of synthetic media, or what is commonly known as deepfakes to support information operations, particularly in the image space. You know, for the past couple of years now, we've been tracking activity sets coming

from pro-China networks that we've tracked using deepfakes. We've tracked pro-Cuban networks. We've tracked a regional network in Argentina that was promoting regional government policies using fake personas, use deepfake images. We've seen the use of deepfakes from Iran and also from numerous non-state actors for political purposes, as well. So, this is another real kind of growing trend in this space.

And then I guess kind of one of the other key developments we've tracked is the extent to which actors are increasingly leveraging kind of the credibility and audiences of legitimate media to help support amplification of their operations. This is what I like to call the IO media nexus, whereby actors will either seek to use the legitimate media to disseminate and amplify, or will seek to mimic legitimate media in some capacity.

So for example, part of the Iran activity set we uncovered in 2018 involved the use of numerous kind of fabricated news sites that were stood up to present as being independent news entities, but were actually being operated by the Iranian state. We see Iran impersonate journalists – both real journalists or kind of stand-up fabricated journalist personas. Again to try and kind of leverage the credibility that journalists have as a source for information. We see the spoofing of legitimate media web outlets, websites, and social media accounts. Again, as an attempt to leverage the existing audiences of these outlets, and kind of present themselves has having more credibility than they actually do.

Luke McNamara:

Well, I definitely want to get into some of the discussion around the IO media nexus and talk about some of the Ghostwriter stuff. But I want to touch on a little bit of what you talked about with the synthetic images, and media, and the deepfake technology. Because that's certainly been an area where there's been a lot of discussion, I guess, for a while now about how that can be used for exactly the sort of threat activity.

ut it's interesting, it seems like from what you're saying a lot of what we've observed to date, while the speculation has been obviously that these sort of technologies can be used for multimedia and video sort of content, it seems like a lot of what we're seeing is just that being utilized for still imagery. Would that be fair? And what are some of the examples we've seen of this?

Lee Foster:

That would be fair. But I'll go even further than that and kind of agree that too much of the public conversation regarding deepfake technology has revolved specifically around the use of deepfake videos. And people always kind of jump to the worst-case scenarios of some fabricated video of a key politician doing or saying something that completely undermines people's trust in the government, as an example. But more realistically, what we're seeing is the use of these kind of technologies to support existing tactics and kind of operational approaches.

So obviously, one very common aspect of many online influence campaigns involves the use of networks of inauthentic social media accounts, right? Pretending to be people belonging to a particular community, maybe from a particular geography, espousing certain political views. A famous example of this, of course, is the Internet Research Agency, which has received so much tension since Russia's 2016 election interference.

And it's areas like this where deepfake technology really has the ability to readily augment existing tactics and techniques. And this is exactly what we see, as you alluded to, with the use of synthetically-generated still images. And these are being used to effectively create avatars, profile photos, for these inauthentic social media accounts. So that makes it harder for people to actually identify these networks as being fake, when they can't, for example, reverse image search a profile picture being used and see that it actually belongs to this real individual who goes by a different name on another social media platform.

So we already see that extensively. And for the most part, it's not technically difficult to do. There are existing tools out there publicly that people can use. In fact, me and a colleague of mine, Phil Tully, recently gave a talk at Black Cat – and listeners can go to our blog – demonstrating just how easy it is for these types of techniques to be used; not only in the image space, but also in the text generation space, and also in the audio generation space.

And these are a couple of areas where I'm increasingly concerned about potential misuse of these types of technologies. Because if you're looking to kind of create content at scale and do it in a way that makes it difficult to attribute it to you, these kinds of tools are extremely useful in that regard.

Luke McNamara:

So it seems like knowing this, knowing that you have a proliferation of newer actors that are entering the space, as well, a lot of different adversary motivations that are here at play from an attribution standpoint – and obviously without getting into discourses of methods, but – it seems like this is becoming a much more complex space that pre-2016, even, to assess the threat activity that's out there. You find the network and campaign.

So I notice in your reporting you're very careful to, where you are making attribution calls, to be very careful about and specific about what you're talking about and saying. But how do you see this space kind of evolving from how you approach this from an attribution standpoint?

Lee Foster:

I think increasingly – and this just doesn't just apply to us, but to really any researchers in this space. I think increasingly we're going to have to start making kind of behavioral assessments that something constitutes an IO campaign. And we can kind of allude to, say, whose political interests it aligns to to give kind of some indication as to why we think an operator might be engaging in this activity. And that's really important for audiences to know. But especially with some of those other trends that I mentioned.

For example, outsourcing to private entities. The attribution to, say, a specific state actor is going to get more and more difficult as time goes on. And so I think we do have to start getting comfortable with the idea of calling out IO activity on a behavioral basis, demonstrating that something is fake, for example. Right? That it's clearly a concerted effort to kind of underhandedly manipulate a particular audience or community without always necessarily being able to get to the specific individual organizations behind that.

Luke McNamara:

So transitioning here to jump into what we've talked about sort of a little bit around at the outset of this, but Ghostwriter. Right? So I think they're a group that personifies at least some of these trends and some of the activity, certainly one of the notable campaigns that your team has uncovered, at least publicly this year. What is Ghostwriter? What have we seen them do? Who are they?

Lee Foster:

So we actually refer to Ghostwriter as an activity set, rather than a group, per se. Because it's likely, actually, that there's overlap in terms of individual actors behind Ghostwriter activity having been involved in other IO campaigns or traditional cyber threat activity. And so we describe this as an activity set based on kind of closely-aligned behavioral overlaps between different kind of sets of incidents.

So what is Ghostwriter? It is a campaign that has leveraged – though not always, but has leveraged – traditional cyber threat activity to promote anti-NATO narratives in Poland, Latvia, and Lithuania, in particular. And kind of the narratives this campaign has been pushing is closely aligned with Russian security interests. It's designed to undermine kind of the image of NATO in those countries to cause local audiences to question the relevance of the organization, or to even see it as a danger.

So what does this activity involve? There's no real kind of generic operation, per se. Each incident seems to vary in ways, but there are a lot of overlaps between them. They will start with the generic creation of some sort of fabricated narrative regarding NATO. This may play upon real-world incidents and kind of manipulate those, and misrepresent those in order to present NATO in a bad light. And it may involve, also, the creation of fabricated material to support that narrative.

For instance, we've seen the fabrication of correspondence coming from NATO officials. We've seen a fabrication of interview transcripts with military officials in those countries. And these are used to then seed fabricated news articles online to spread these narratives. And then this is where kind of the mix of dissemination tactics emerges. In many of these cases, there's been the compromise of legitimate media outlets, or other websites, any stimuli. And then those fabricated articles have been posted to those sites.

We go back to what I was talking about earlier about this IO media nexus idea, it's a classic example there, right? The operators are leveraging a legitimate audience and perceived credibility of those media sites to push this fabricated content out there. In other instances when we've seen kind of operator spoof e-mail accounts belonging to government or military officials, and then e-mail those narratives, those articles, to legitimate media outlets to try and get them to pick them up then.

On top of that, we've seen the extensive use of false personas, often presenting as journalists or analysts in those countries, publishing these self-written articles to numerous websites that allow for user-generated content – in English language, predominantly. And if you go to the report – which is, we have a blog on this activity and a report linked to it – you can see there some of the core sites that they've used in this activity.

And then we've also seen these types of articles being pushed by, pushed on various blogs and web pages that we believe are actually Ghostwriter-controlled. So really just a mix of dissemination tactics this campaign uses.

Luke McNamara:

And this is activity that we track separately from another set of activity that I guess a lot of people would be familiar with, secondary infection. What are some of the characteristics of how this – and you described a lot of the details of the activity itself. But how does it differ from what we've seen from the secondary infection networks and activity?

Lee Foster:

So there are numerous differences between the two activity sets. But a couple of the kind of key ones really would be firstly just the level of use of traditional cyber threat activity, i.e. intrusions or e-mail spoofing that we've seen from the Ghostwriter campaign. Currently, there's only one potential example of secondary infection activity being supported by traditional and cyber threat activity.

And that is with regards to last year's UK trade documents prior to British elections. However, the sourcing for that currently is just anonymous government officials in the UK. But even if that does turn out to be accurate, that's the only case thus far we've seen of kind of traditional cyber threat activity being used to support secondary infection activity. Whereas it's really been quite a prominent characteristic. Although, as I mentioned earlier, not universal for Ghostwriter activity.

Another key difference is the nature of the false personas that are used in the two campaigns. Secondary infection, one of its kind of key characteristics is the use of single-use burner accounts on various platforms online, whereby when a narrative is spun up or a piece of fabricated content is spun up and disseminated, these single-use accounts are used for them.

They post it to a user, a content-generated site. You know, to push it out. And that's it. That account is never used again. And that's really kind of been the case for all secondary infection operations.

In the Ghostwriter context, many of the personas that Ghostwriter has used are multi-use and have kind of deeper histories established. As mentioned before, they actually present themselves as being journalists, or analysts, and so on. And in a couple of cases they've actually stolen the identity of real individuals. For example, impersonating a real journalist. And these personas have a more extensive history of activity. They've posted numerous articles over time to these various sites.

Luke McNamara:

So when you're thinking about the significance of this activity from a capability standpoint, from a messaging and targeting standpoint, how does this rank compared to other threat activity that we've seen in the past? Obviously, one of the most difficult things about this if we're talking significance is what has been the impact to the communities that they're trying to message to? And I know that's a very difficult question to answer. But in looking at kind of all of this, how would you rank the significance of this activity set?

Lee Foster:

I would look at it as more significant than people might look at it on face value, just because it's a narrow activity set in terms of its narrative targeting, you know, of these specific countries. And people might dismiss it because of that. But I think it's really important to remember we think about kind of Russia-aligned and Russian-attributed activity, in particular. Historically, we've long-seen kind of a transition of activity from Eastern Europe to elsewhere. Right?

And so I think it's important that we don't kind of neglect this activity set just because it so happens to be specifically anti-NATO, specifically in these three countries right now. All of the tactics that I just mentioned are readily deployable elsewhere. So we should consider it significant in that sense, because it shows that even today there is this close marriage between traditional cyber threat intrusion activity and online-driven influence operations that have kind of come to dominate kind of the geopolitical conversation over the past few years.

Luke McNamara:

And for activity sets like that that do incorporate some more traditional intrusion activity that we typically see from a lot of the nation state APT groups that we track, does that kind of afford us an additional way to combine different methods of doing attribution and looking at the totality of this?

Lee Foster:

Yeah. Absolutely. When, obviously, you start having cyber threat activity, the core part of these campaigns, that is going to leave behind some forensic evidence that people will be able to use to aid in attribution. And so having the ability to merge both the open source IO analysis and investigatory kind of skill set with the traditional kind of cyber forensics skill set that really does allow more avenues for attribution to open up there.

Luke McNamara:

So transitioning to what we might see next. And obviously, part of this question is going to be tied to US elections and disinformation, since that is a topic of universal interest right now is what we might see in November, particularly from threat actors engaged in disinformation. For people that are watching this space carefully with what we've seen employed by activity sets like Ghostwriter, secondary infection, and from other threat actors elsewhere, what should they be looking for that we might see? What might be some things to keep an eye out for when it comes to disinformation in the 2020 US elections?

Lee Foster:

So the first thing I would highlight is just the continuation of some of those trends that I kind of outlined at the beginning of our talk here. Right? I think we're going to see increasing use of outsourcing by state actors behind these campaigns. We're going to observe increasing interest in actors of using synthetic media to support their operations. Not necessarily in the lead-up to the elections, but kind of long-term. We're going to see more and more players start experimenting with these kind of tactics for their own political purposes. And in particular, I think we're going to start seeing more and more non-state actors try to utilize some of these capabilities. So that's definitely something that I would be kind of on the lookout for.

As it pertains to the election specifically, again, I think we have to be conscious of domestic actors in this space, and make sure we're paying attention to that, and not just kind of looking

for Russia, which is still kind of just the, for some people it's all they kind of picture when they think of online influence operations. Right? And it's just as I mentioned before, there's so many more actors involved in this space. And they all have different interests as it pertains to US politics. So we have to keep an eye out for those.

I think the one thing we're likely to see is actually immediately post-election, will be various actors trying to call into question the legitimacy of the results. And this, I think, will be the case irrespective of who actually wins the election. Because different state actors and different non-state actors, they have different agendas at play. They have different preferences as it pertains to the outcome of the November elections. So I would anticipate quickly seeing kind of efforts to delegitimize the result, whatever that may be.

Luke McNamara:

What about from the standpoint of if we look at some of the TTPs of Ghostwriter, specifically the targeting of local media outlets, and you think about some of those smaller media outlets that may not have significant security budgets, and the resources to really defend their networks or their infrastructure, their websites that can be hijacked and used in those means; how do you think that sort of tactic could be employed here? Do you think it would have the same sort of efficacy? Or is the media landscape environment very different from what we saw in parts of Europe where this activity was present?

Lee Foster:

So that would all really depend on who was being targeted and how. Right? I do worry about kind of an eleventh-hour, say, compromise of a prominent news entity that is used to push out some sort of fabricated article. And it's quickly retracted, and so on. But by then it's kind of too late. Right? This kind of politics moves so quickly and people will just incorporate *[laughs]* that narrative into their beliefs. So that worries me.

But you suddenly hit on something important, I think, as it pertains to the susceptibility of smaller, more local news outlets. We've seen this, in fact, with some activity we previously assessed to be of Iranian origin, a campaign we call Distinguished Impersonator, by letters that were being submitted by these operators to local papers in the US that were then publishing. And these tended to be critical of the current administration and of kind of Iran's rivals in the Middle East.

But it shows you how susceptible some entities can be to this type of activity. And that didn't even involve any kind of compromise of back-end systems to publish these.

Luke McNamara:

Yeah. A very scary thing to consider, but realistic in terms of how we've at least seen this employed elsewhere. So I guess in wrapping this up. And obviously, we've spent time talking about a lot of different interesting trends that you're seeing around information operations today. What might be some predictions moving away from the US elections of activity that we see going forward in the coming year?

Obviously, we talked about how some things that we've expected to see or had been called for, deepfakes involving videos and multimedia, we've not quite seen that yet. What might be some things that maybe we're not considering now? Or even some regions and areas where we might see increased information operations, given that this capability seems like it's of value and interest to a lot of emerging threat actors.

Lee Foster:

I think we're likely to see just greater adoption of these types of tactics, in general, by actors around the world. And it's important to acknowledge that this isn't always of the nature that the people here tend to conceptualize this threat as X country targeting the US voters.

A lot of the activity we do track is domestic in nature. Right? It's governments targeting their own people with particular narratives and messaging. Or it's opposition groups targeting their own people with anti-government narratives. Or it plays out on kind of a smaller regional scale. It's different neighbors in the Middle East or Latin America and kind of targeting audiences in each other's countries.

So, that's something I would envision we continue to see grow as these kind of techniques get a broader exposure, and various actors look to see perceived benefits from these types of campaigns. And I think the same is true for non-state actors, as well, in general. We've already tracked various incidents involving specific individuals abroad targeting US politics before. There's any number of motivations that might prompt someone to do this.

But again, as I highlighted previously, these types of techniques are so readily employable and barriers to entry so low, particularly when they don't use aspects of traditional cyber threat activity like intrusions, that anyone with a particular agenda can reasonably kind of develop a campaign. So that's hugely concerning.

Luke McNamara:

Well, that's a very cheery closing thought to kind of leave the audience with. To the point of increasing numbers of actors already entering the space, and we'll continue to probably see that, I think it's great work that you and your team are doing to keep us educated as to what the threat landscape looks like. If people are more interested in learning about some of the campaigns discussed here, Ghostwriter, where can they go to find out that information you mentioned earlier?

Lee Foster:

So if you just go to the FireEye Threat blog, there will be a blog there on Ghostwriter. There will also be a blog on the use of synthetic media and IO that I mentioned, as well, which is a more technical read, but really does try to outline that the technical sophistication that people believe is required for the use of these types of tools is not that high. And so it is something that we need to keep an eye on.

Luke McNamara:

Awesome. Well, definitely check those resources out. And Lee, thanks again for coming on and talking to us about IO.

Lee Foster:

Great. Thanks for having me.