A Last Clever Knot?

horkos.medium.com/a-last-clever-knot-26fd26765e8d

Horkos November 24, 2020

Horkos

Oct 23, 2020

.

7 min read

“It’s the start, I should tell you now, of a very, very, clever knot.” (from the BBC’s Tinker Tailor Soldier Spy)

Note: This blog has been updated as of 19 November 2020 to incorporate additional material, which is preceded by an asterisk.

On 22 October, 2020, the U.S. government issued an alert (AA20–296A) for a widespread campaign of malicious activity by a Russian state-sponsored cyber operations entity tracked under a variety of industry monikers. This activity took place between at least September 2020 and October 2020, with its targeting focused on a large volume of entities associated with U.S. state, local, tribal, and territorial (SLTT) networks and aviation sector networks. The Russian entity named as responsible has been linked to Russia's Federal Security Service (FSB) by the Washington Post in connection to past activity attributed by the U.S. government to "Russian government cyber actors" and by the New York Times in relation to AA20–296A; for the sake of argument let's say that attribution is plausible, if not correct. (*For additional context, Joe Slowik has also published an exceptional blog on the history and implications of this entity and its operations.) I'd rather not try to summarize all the activity described in the alert here, but instead to address its operational logic and how it can be taken as a likely exemplar of Russian active measures: an insidious blend of offensive counterintelligence and influence operations that, in the words of George Smiley, "would be beautiful in another context."

The operation's apparent targeting parameters provide our first clear clue. While the scope was clearly focused geographically on the United States and sectorally on SLTT-related entities, the described scale — in both the text of the alert (e.g., "wide variety of U.S. targets", "targeted dozens of SLTT government and aviation networks", a one-off reference to at least one apparent education sector target likely related to SLTT networks) and a *subsequently published heat map — is relatively large. To my eye, an operation so large against that target scope casts some doubt on as to how much actual value its operators reasonably could have expected to draw from the activity if the intention was intelligence collection in advance of the U.S. presidential election in November 2020. Regardless of if it relied on the same logic, the U.S. government seems to have arrived at a similar conclusion, not mentioning a possible collection motive but rather that "…the actor may be seeking access to obtain future disruption options, to influence U.S. policies and actions, or to delegitimize SLTT government entities."

I would take that conclusion a bit further. I believe these targeting parameters suggest a different motive than either intelligence collection or acquiring network/system disruption options, because in either case — facing the election day time crunch — the actor would focus their targeting much more narrowly if they really needed to deliver results. I am disinclined to believe the FSB (for the sake of argument) would invest the time and effort required of such a large operation, so close to a major event of strategic significance (i.e., the U.S. presidential election), if it was really meant as a shotgun approach to acquire options or data. If those requirements existed, they would've certainly predated September 2020 and been the subject of more highly-targeted activity before autumn.

Our second clue is the tradecraft employed: exploit-chaining that relies on the "CVE greatest hits" of the last few years. Given the aforementioned asymmetries within the targeting parameters, the picture I see emerging is something akin to the actor artificially constructing

a game reserve and then opportunistically hunting within it. The targeting parameters are the game reserve: entities in the United States, associated with SLTT government no matter how tangentially. The heavy reliance on exploit chaining — easily enough automated — is the opportunistic hunting, so opportunistic that it borders on indiscriminate. Yet it remains confined to the boundaries of the artificial reserve. With that kind of model, a viable goal emerges in my mind: the offensive counterintelligence effect known as degradation.

I have previously talked about the potential degradation functions of large-scale targeted intrusion campaigns. Such campaigns require defensive efforts to be expended by a large group of particular public and private sector entities. From initial incident responders potentially all the way up to senior government officials, this effort ripples upward and outward as it gathers steam. From a counterintelligence perspective, this exertion and the anxieties that accompany it serve the function of frustrating an adversary on multiple levels. (In an academic sense, it is an interesting example of capitalizing on Clausewitzian friction.) The anxiety element of such friction also presents further exploitable possibilities.

Specifically, an actor could harness the collateral psychological effects of such a degradation-oriented campaign for the purpose of active measures, including influence operations, against a wide variety of audiences. In essence, the psychological stress endured by blue team personnel and decision-makers is both the core degradation effect as well as the first-order influence impact. The second-order influence impacts would be directed at any individuals that responders et al. would tell about their trials and tribulations, including individuals who may inform the press or general public. Thus, the third-order influence impacts could affect at least a portion of the wider public. And once you are trying to get average people to wrap their heads around a campaign like that, you can count on pitfalls of the domestic information environment to help do some of the heavy lifting for the adversary when it comes to FUD (fear, uncertainty, doubt) generation. The anxiety that was primarily designed to wear down security and government communities now has actual potential to cause alarm among a susceptible percentage of the public.

Do you see where I'm going with this?

I personally think there's a good chance that the reported activity was intended primarily as a degradation operation targeting the U.S. government and secondarily as an influence operation against both the U.S. government and the American public. In that sense, I'm advocating for the potential motives of "influence U.S. policies and actions" and "delegitimize SLTT government entities" as described in AA20–296A. I base this on the reported targeting parameters and tradecraft employed, as well as the effects possible from both counterintelligence- and active measures-oriented operational planning perspectives. To return to my earlier metaphor: the opportunistic hunting within the chosen game reserve in such close temporal proximity to the election would present ideal conditions for that ripple-upward-and-outward anxiety that is so ripe for exploitation. In particular, the politicization of

the integrity of the election presents some extraordinary opportunities for undermining faith in democratic institutions — to include elections and their outcomes — and maybe sow a little general chaos in the process. It would be most Chekist in its logic.

There is, again to paraphrase le Carré's Smiley, a last clever knot. I think of it as a dead man's switch for the influence operation goals of this activity. It revolves around whether or not the activity is disclosed and ensuring some level of effects across a variety of scenarios. If the activity is discovered but not publicly disclosed, at least the first-order targets and some of the second-order targets have been impacted. If the activity is made public by entities other than the U.S. government, then first-order, second-order, and third-order target sets are impacted to at least some degree. If the U.S. government chooses to disclose the activity, all three orders of targets are impacted and the potential for general alarm among the public increases because an official government announcement almost guarantees a greater minimum level of public consumption (including via media amplification). It even may be the actor's intent to make this latter potential evident in some way to government decision-makers as an attempt to create a disclination to disclose or attribute the activity, knowing that internal debate over what course of action to take could intensify friction within relevant government entities.

However, that third option — disclosure and attribution by the U.S. government — offers significant positive trade-offs that could counter the some of the intended effects of such activity, if done properly. Being the first to disclose, and disclosing in sufficient detail, would grant the U.S. government a strong opportunity to set the narrative and in doing so potentially engage in some meaningful inoculation of the general public against misunderstanding, manipulation, and panic. It also presents a chance to throw the adversary off their operational timelines, as a bit of retaliatory degradation. There of course remains the very real risk that the actor responsible may still leverage knowledge gained in this activity in a variety of ways both during and after election day (e.g., so-called "perception hacks", hack-and-leak operations, website defacements, system/network disruptions, etc.); however, those possibilities would exist regardless of if the underlying activity was disclosed or not.

*There also exists the risk that domestic actors, not just the foreign ones associated with the SLTT-targeted intrusion activity, could attempt to use the existence of the aforementioned activity to bolster independent efforts (e.g., legal challenges, influence narratives, etc.) meant call into question the validity of the outcome of the election. Kudos to Kyle Ehmke for pointing out this very important point. It is likely that the actor responsible for the activity covered in AA20–296A anticipated, and even desired, that their actions could be potentially exploited in such a manner by domestic actors.

In the spirt of that third option well-executed, I believe the 22 October disclosure represents a meaningful response to Russian active measures. Some may view it as a pyrrhic victory in the absence of some kind of tangible cost imposition like indictments or disruption of adversary networks, but to take that stance ignores the more intangible, subtle dynamics likely at play. To be clear — I do not attempt to say the theory I've outlined definitely is what

this activity represents, merely that I think there is sufficient evidence to make a good case for it. In the same vein, I must acknowledge there are details about the reported activity that I do not know and those details could weaken or outright contradict my theory. However, at the end of the day, I fundamentally remain impressed with the adversary's apparent logic — at least from how I perceive it — and feel further secured in the belief that to not respect your adversary's capacity for cunning is hubristic.

Disclaimer: As a reminder, all views expressed on this blog, including this post, solely represent my personal views and not those my employer.