A XENOTIME to Remember: Veles in the Wild

Published: 2019-04-12 · Archived: 2026-04-06 02:06:30 UTC

"When I use a word," Humpty Dumpty said, in rather a scornful tone, "it means just what I choose it to mean—neither more nor less." – *Through the Looking Glass*, Lewis Carroll

FireEye recently published a blog covering the tactics, techniques, and procedures (TTPs) for the "TRITON actor" when preparing to deploy the TRITON/TRISIS malware framework in 2017. Overall, the post does a commendable job in making public findings previously only privately shared (presumably by FireEye, and in several reports I authored for my employer, Dragos) to threat intelligence customers. As such, the blog continues to push forward the narrative of how ICS attacks are enabled through prepositioning and initial intrusion operations – an item I have discussed at length.

Yet one point of confusion in the blog comes at the very start: referring to the entity responsible for TRITON as the "TRITON actor". This seems confusing as FireEye earlier publicly declared the "TRITON actor" as a discrete entity, linked to a Russian research institution, and christened it as "TEMP.Veles". In the 2018 public posting announcing TEMP.Veles, FireEye researchers noted that the institute in question at least supported TEMP.Veles activity in deploying TRITON, with subsequent public presentations at Cyberwarcon and the Kaspersky Lab-sponsored Security Analyst Summit essentially linking TRITON and the research institute (and therefore TEMP.Veles) as one in the same. Yet the most-recent posting covering TTPs from initial access through prerequisites to enable final delivery of effects on target (deploying TRITON/TRISIS) avoids the use of the TEMP.Veles term entirely. In subsequent discussion, FireEye personnel indicate that there was not "an avalanche of evidence to substantiate" anything more than "TRITON actor" – summing matters by indicating this term "is the best we've got for the public for now".*

Meanwhile, parallel work at Dragos (my employer, where I have performed significant work on the activity described above) uncovered similar conclusions concerning TTPs and behaviors, for both the 2017 event and subsequent activity in other industrial sectors. Utilizing Diamond Model methodology for characterizing activity by behaviors attached to victims, we began tracking TRITON/TRISIS and immediate enabling activity as a distinct activity group (collection of behaviors, infrastructure, and victimology) designated XENOTIME. Based on information gained from discussion with the initial TRITON/TRISIS responders and subsequent work on follow-on activity by this entity, Dragos developed a comprehensive (public) picture of adversary activity roughly matching FireEye's analysis published in April 2019, described in various media.

At this stage, we have two similar, parallel constructions of events – the *how* behind the immediate deployment and execution of TRITON/TRISIS – yet dramatically different responses in terms of attribution and labeling. Since late 2018, based upon the most-recent posting, FireEye appears to have "walked back" the previously-used terminology of TEMP.Veles and instead refers rather cryptically to the "TRITON actor", while Dragos leveraged identified behaviors to consistently refer to an activity group, XENOTIME. Given that both organizations appear to describe similar (if not identical) activity, any reasonable person could (and should) ask – why the inconsistency in naming and identification?

Aside from the competitive vendor naming landscape (which I am *not* a fan of in cases on direct overlap, but which has more to say for itself when different methodologies are employed around similar observations), the distinction between FireEye and Dragos' approaches with respect to the "TRITON actor" comes down to fundamental philosophical differences in methodology. As wonderfully described in a recent public posting, FireEye adheres to a naming convention based upon extensive data collection and activity comparison, designed to yield the identification of a discrete, identifiable entity responsible for a given collection of activity. This technique is precise and praiseworthy – yet at the same time, appears so rigorous as to impose limitations on the ability to dynamically adjust and adapt to emerging adversary activity. (Or for that matter, even categorize otherwise well-known historical actors operating to the present day, such as Turla.)

FireEye's methodology may have particular limitations in instances where adversaries (such as XENOTIME and presumably TEMP.Veles) rely upon extensive use of publicly-available, commonly-used tools with limited amounts of customization. In such cases, utilizing purely technical approaches for differentiation (an issue I lightly touched on in a recent post) becomes problematic, especially when trying to define attribution to specific, "who-based" entities (such as a Russian research institute). My understanding is FireEye labels entities where definitive attribution is not yet possible with the "TEMP" moniker (hence, TEMP.Veles) – yet in this case FireEye developed and deployed the label, then appeared to move away from it in subsequent reporting. Based on the public blog post – which also indicated that FireEye is responding to an intrusion at a second facility featuring the same or similar observations – this is presumably not for lack of evidence, yet the "downgrade" occurs all the same.

In comparison, XENOTIME was defined based on principles of infrastructure (compromised third-party infrastructure and various networks associated with several Russian research institutions), capabilities (publicly- and commercially-available tools with varying levels of customization) and targeting (an issue not meant for discussion in this blog). In personally responding to several incidents across multiple industry sectors since early 2018 matching TTPs from the TRITON/TRISIS event, these items proved consistent and supported the creation of the XENOTIME activity group. This naming decision was founded upon the underlying methodology described in the Diamond Model of intrusion analysis. As such, this decision does not necessarily refer to a specific institution, but rather a collection of observations and behaviors observed across multiple, similarly-situated victims. Of note, this methodology of naming abstracts away the "who" element – XENOTIME may represent a single discrete entity (such as a Russian research institution) or several entities working in coordination in a roughly repeatable, similar manner across multiple events. Ultimately, the epistemic foundation of the behavior-based naming approach makes this irrelevant for tracking (and labeling for convenience sake) observations.

Much like the observers watching the shadows of objects cast upon the wall of the cave, these two definitions (XENOTIME and TEMP.Veles, both presumably referring to "the TRITON actor") describe the same phenomena, yet at the same time appear different. This question of perception and accuracy rests upon the underlying epistemic framework and the goal conceived for that framework in defining an adversary: FireEye's methodology follows a deductive approach requiring the collection of significant evidence over time to yield a conclusion that will be necessary given the premises (the totality of evidence suggests APTxx); the Dragos approach instead seeks an inductive approach, where premises may all be true but the conclusion need not necessarily follow from them given changes in premises over time or other observations not contained within the set (thus, identified behaviors strongly suggests an activity group, defined as X).

From an external analysts' point of view, the wonder is, which is superior to the other? And my answer for this is: neither is perfect, but both are useful – depending upon your goals and objectives. But rather than trying to pursue some comparison between the two for identification of superiority (an approach that will result in unproductive argument and social media warring), the point of this post is to highlight the distinctions between these approaches and how – in the case of "the TRITON actor" – they result in noticeably different conclusions from similar datasets.

One reason for the distinction may be differences in evidence, as FireEye's public reporting notes two distinct events of which they are aware of and have responded to related to "the TRITON actor" while Dragos has been engaged several instances – thus, Dragos would possess more evidence to cement the definition of an activity group, while FireEye's data collection-centric approach would require far more observations to yield an "APT". Yet irrespective of this, it is confusing why the previously-declared "TEMP" category was walked back as this has led to not small amount of confusion – in both technical and non-technical audiences – as to just what FireEye's blog post refers.

Thus respected journalists (at least by me) conflate the "TRITON actor is active at another site" with "TRITON malware was identified at another site". In this case, we're seeing a definite problem with the overly-conservative naming approach used as it engenders confusion in a significant subset of the intended audience. While some may dismiss adversary or activity naming as so much marketing, having a distinct label for something allows for clearer communication and more accurate discussion. Furthermore, conflating adversaries with tools, since tools can be repurposed or used by other entities than those first observed deploying them, leads to further potential confusion as the "X actor" is quickly compressed in the minds of some to refer to any and all instantiations of tool "X".

Overall, the discussion above may appear so much splitting of hairs or determining how many angels can dance on the head of a pin – yet given the communicative impacts behind different naming and labeling conventions, this exploration seems not merely useful but necessary. Understanding the "how" and "why" behind different entity classifications of similar (or even the same) activity allows us to move beyond the dismissive approach of "everyone has their names for marketing purposes" to a more productive mindset that grasps the fundamental methodologies that (should) drive these decisions.

*Note: Following publication, John Hultquist, director of FireEye iSight, provided clarification on use of "TEMP" naming criteria in FireEye public reporting via Twitter.
