Dear Chandan,

Thank you for your detailed response. While you raised many topics, we believe that our analysis is valid. Please see below -

You wrote:
First, we observed that the testing conditions used by your research team departed substantially from an appropriate real-world implementation of Prisma Cloud. Instead, they focus solely on a narrow set of runtime defense features, while completely ignoring other features of Prisma Cloud, some of which are implemented in the development stage and which prevent real-world runtime risks.

Our response:
In our analysis we reviewed the full set of runtime defense features of Prisma cloud, which your product management team described as “robust runtime monitoring and protection mechanism”. We totally agree that the best prevention is much earlier than runtime, but the purpose of this analysis was to review the proclaimed robust runtime monitoring and protection mechanism - Which came up extremely short.

You wrote:
Second, perhaps because Orca’s research team only conducted a “quick review”, you may not have been aware that the analysis was apparently conducted on Prisma Cloud Compute version 20.09, an older product version released last year that is no longer supported

Our response:
Prisma cloud compute version 20.09 was released in September-2020 - just a few months before the test was conducted. According to your support policy, the current and previous versions are supported - quoting “Currently, the Compute support lifecycle is the current major version and the previous major version ("N-1").”. Naturally, only features available in this version at this time were evaluated.
We would love to re-run the tests with an updated version - for the industry sake - will you be willing to provide us with access to an updated version with the new features and approval to use it for the assessment that will be later published?

You wrote:
Third, you refer to cryptominer and lateral movement detection which, your research team asserts, can be thwarted by simple changes to file names. This fails to evaluate the outcome of such file name changes in a real-world context. Filenames are merely one of many dimensions Prisma observes to categorize runtime risk and, if the file does not perform any malicious activity such as vertical or horizontal port scanning, there will be no further signals to identify potential risk. Had the renamed executable been actually used to perform a malicious activity, the system would have detected it and, depending on a customer’s configuration, prevented it.

Our response:
Can you elaborate on what you mean 'the outcome of such file name changes in a real-world context'? In the example we showed in the video, we renamed the cryptominer to a different name - and it went undetected although its execution wasn't affected by the changed name. It is unclear why a cryptominer should perform 'malicious activities such as vertical or horizontal port scanning'. In that case, the cryptominer performed as intended (that was supposed to be detected by the platform), and went undetected simply as the name was changed.

You wrote :
Fifth, you assert that "[m]alware detection is extremely limited - not detecting even extremely well-known malware." We disagree. The examples used in your video are not real malware but rather dummy binaries used for academic tests. Prisma does not have the dummy files in its database in order to avoid triggering false positives. Here again, your team's focus on the out-of-date Prisma product also resulted in your testing against outdated antimalware capabilities, which did not include integration with industry leading WildFire services.

Our response:
Eicar is a standard test file any anti-malware tool should detect. Having said that, we tested 3 files in the video - Eicar as well as 2 well known linux malware - turla and xbash. Turla and Xbash have been well known malware in the industry for a long time.

You wrote:
Finally, you claim that "[e]xploit detection doesn't work - even for well-known exploits that are supposedly understood by your platform." Again, we disagree. The basis for your assertion appears to be a test not representative of an actual attack. Prisma Cloud detects and prevents real-world threats and actual attack patterns. Whomever is recording the video never downloads or runs additional binaries. The only action taken in the video was to run an 'ls' equivalent within the PHP framework rather than simulating an actual attack, such as downloading additional binaries or even running ones not included in the normal operation of the container.

Our response:
Your documentation states that "Anti-malware provides a set of capabilities that lets you alert or prevent malware activity and exploit attempts.". The example we showed included a successful exploit attempt. Can you elaborate on the real platform capabilities - are you protecting against exploit attempts (as your documentation states) or merely hoping that post exploitation operations will be detected by other means? Are you going to update your customers that rely on the proclaimed exploit detection capabilities?

You wrote:
By using an artificial and unrepresentative testing environment, along with an unsupported version of

Prisma, we believe your research team's quick review could not meaningfully assess the true capabilities of Prisma.

Our response:
As noted earlier, the version was a few months old at the time of testing (just released in September-2020), and according to your documentation is still supported.
At the end of the day, it isn't up to us or Palo Alto to decide whether this testing environment is representative of the real world or not. This is up to the prospective customers who are assessing their cloud security strategy and evaluating different solutions. Therefore, we would love to continue this discussion in writing, as I believe it can provide valuable insights to prospective clients.

Based on your responses, we see that the Palo Alto view is that the above observations aren't issues that you're planning to fix, and as such we are not subject to any responsible disclosure guidelines. Can you please confirm that?

Yours,

Avi Shua