Dear Mr. Shua,

Thank you for your email below regarding the runtime monitoring and protection capabilities of Palo Alto's Prisma Cloud Security suite. I wanted to note a few things in response to your positions.

At the outset, we wish to thank you for bringing these matters to our attention and for what we hope will be your good faith and continued, constructive engagement with us regarding the resolution of the identified issues. As a longstanding member of the security community, Palo Alto Networks welcomes third-party reports of potential vulnerabilities in its products and services, and it has a well-developed process in which it engages its Product Security Incident Response Team (PSIRT) to investigate and respond promptly to relevant reports. Pursuant to this process, we invite your research team to present and discuss its observations with the Prisma Cloud security experts. Please let us know and we can arrange such a meeting at your earliest convenience.

Nevertheless, we would like to take this opportunity to respond to a few of the issues you have raised. 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.

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 and which does not have Wildfire Integration, which safely evaluates unknown binaries before allowing them to run. Further, the Intelligence Stream for the Prisma Cloud version your team ran had been disconnected for nearly two months, meaning that threat data (IP reputation lists, DNS block lists, malware signatures) was out of date. We would hope that your team would focus on current and fully supported versions in the future.

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.

Fourth, your research team's review purports to find that "C&C communication detection is evaded when communicating using IP addresses." As you know, no static data sets of malicious IPs and DNS names will cover every single malicious endpoint on the internet. They constantly change. This is why the failure to have Prisma Cloud's Intelligence Stream up-to-date and enabled was a particular problem in this case. The Intelligence Stream draws from malicious

endpoint data from AutoFocus, Palo Alto Networks' global database of intelligence sourced from Unit 42 researchers and thousands of sensors globally. That said, static threat feeds are just a single layer of the Prisma runtime defense. Other mechanisms, for example, automatically detect the installation and execution of the Tor binaries and their outbound egress ports as anomalous.

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.
While static malware detection is one layer of Prisma's process runtime defense, most process runtime defense capability is provided by our ML-generated models that act as explicit "allow" lists. Malware signatures and WildFire integration provide additional flexibility in determining response actions when unknown binaries are encountered, but Prisma Cloud is not reliant upon them exclusively. By default, an unknown binary that is not part of a container's original, genuine image is prevented from running regardless of whether any static signature exists.

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. 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.
We sincerely hope that you will accept our invitation to engage directly with the PSIRT to work through these issues in detail. While we are confident in our understanding of the product, if we are incorrect in this response, we invite you to show us why when we meet.

Regards,
Chandan