InfoWorld

JANUARY 31, 2018

GET TECHNOLOGY RIGHT®

INSIDER

Review: BluVector enables machines to protect themselves

With machine learning that gets smarter and more network-aware over time, BluVector can tip the scales back in favor of defenders.

Network security programs and the human IT operators who manage them are under constant threat. New attack techniques like malware deployed without files are straining resources and testing defenses in two critical ways. First, brand new threats and attack techniques often have at least a small window of time when they can bypass some defenses before defenders catch up. Second, even if critical threats like zero-day malware are stopped, the constant siege of attackers means that defenders are likely to get overloaded by both real alerts and false positives.

One possible solution that has only recently become an option is tasking machines and computers with protecting themselves. If a security program could be programmed to think and act like an analyst, then it could try and counter malware and human-backed intrusions at machine speed, a move that would give defenders a serious home court advantage.

The BluVector defense does just that, offering advanced detection and response, and even threat hunting, all performed at machine speeds. BluVector works almost right away, but also has deep machine learning capabilities, so it gets even smarter over time. It will learn the intricacies of each network that deploys it, tweaking its algorithms and detection engines in a way that makes the most sense for the environment.

BluVector is installed as either a hardware-based network appliance, or as a virtual machine. It can operate in-line with network traffic, stopping and remediating threats in real time as they attempt to enter a protected space, or as a retrospective tool that can scan the work performed by other programs and analysts, catching threats that they might have missed and recommending fixes. It is designed to work with all IPv6 traffic as well as older IPv4 streams, so it can operate in environments that are rich in internet of things (IoT) and supervisory control and data acquisition (SCADA) devices, such as those in industrial and manufacturing settings, as well as for normal office type environments.

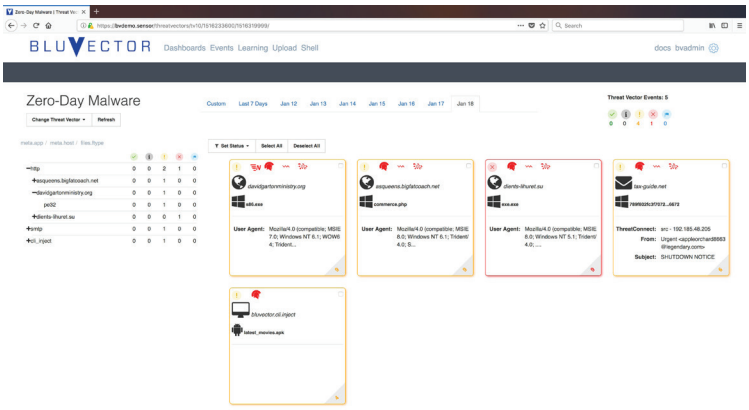

John Breeden II/IDG Captured or suspected threats appear as cards in the dashboard showing every detection engine that flagged it, and which ones let it pass. This is the first stop in identifying or confirming a zero-day threat.

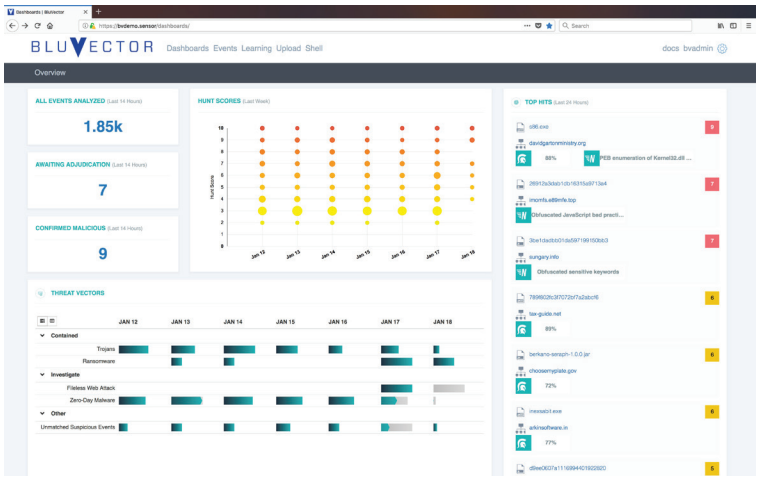

John Breeden II/IDG The main dashboard for BluVector provides a jumping off point for threat hunters to investigate incidents caught by the suite of detection engines. There is a lot of valuable information here, but it can be completely ignored if BluVector is allowed to operate independently without restrictions.

It would be extremely easy for BluVector to operate like a black box, performing its magic behind the scenes without really explaining how it all works. The engineers who designed the system took a different path however, letting humans see everything the system is doing, and also why it is taking the actions it does, or, in the case of offline deployments, why it recommends a particular course of action.

Broken down into components, the BluVector program is basically a series of detection engines that in turn feeds into a probability engine that determines what actions to take. From there, it can automate responses quickly enough to keep advanced threats out of a network, or alert human users for authorizations – a process that keeps humans in the loop, but slows down protection and eliminates one of BluVector's core advantages. Still, organizations wary about turning over cybersecurity functions to machines have the option to keep humans informed, or to keep their hands on the tiller.

The BluVector detection engines tap into supervised machine learning, speculative code execution, behavioral heuristics, signatures and threat intelligence rules, a file extractor and a portable executable scanner, each with its own engine. This feeds data into the probability engine, which assigns a score to each flagged file or piece of code. BluVector can find code inside of scripts or within traffic streams, and pull it out to be reassembled and run through the detection engines.

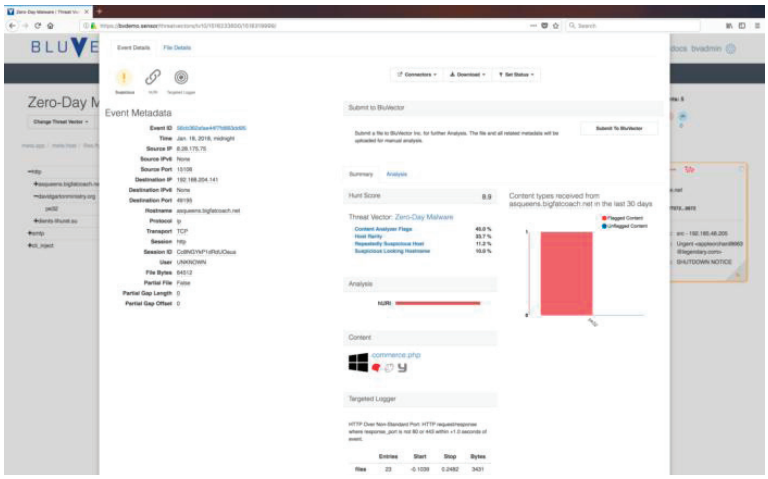

John Breeden II/IDG Code or programs flagged as malicious can be examined within the BluVector interface, which breaks down all the reasoning for the threat designation.

Basically, the BluVector box is performing threat hunting using artificial intelligence (AI), very much like a human analyst would given the time, training and resources. Using threat hunting to find unknown threats is a highly useful skill in cybersecurity these days, with a critical shortage of IT professionals who practice that discipline. This has led to companies offering threat hunting as a service. BluVector takes that concept even further, by having an intelligent machine handle the hunting.

During our testing, we observed how BluVector was able to catch some threats that more traditional defenses like antivirus and signatures missed. Diving into the detection explanation, it was clear that the file in question was caught by the internally developed threat engine called Hector, while missed by others such as antivirus. Looking closer, we could see exactly what factors contributed to the program flagging what was possibility a zero-day threat. This included concerns over what the code wanted to execute, the rarity of the host sending the file and other factors like a strange nomenclature that would have probably tipped off a human threat hunter had they been looking in the right place at the right time.

The program then helpfully collected every piece of information that a threat hunter would find relevant for an investigation, such as HTTP headers, associated files, DNS logs, histories and user data and protocols. All of this would make BluVector an excellent tool for a human threat hunter. The difference is that, if set to operate in real-time, BluVector would have already performed the investigation and remediated the situation before a human even had their first cup of coffee. It can spring into action nearly instantly at the first sign of trouble, so it can catch threats and eliminate them at the network perimeter, a feat very few defenses can accomplish anymore.

Where Hector works on files, the Nema detection engine works on code and snippets of code. In our tests, BluVector was able to extract a piece of file-less malware for examination from the overall stream, and properly identify it as such. Nema identified it as malware based on several factors such as terrible programming practices, suspicious string entropy, obfuscated URL protocols and others. Thereafter, it worked just like a Hector detection, explaining everything that went into its decision to classify the threat as a previously unknown malware, and collecting all supporting data needed for a threat hunter, analyst or auditor. And, if allowed to do so, BluVector could stop that advanced threat before it could breach the network perimeter.

The OpenAPI nature of BluVector means that it can easily integrate into an existing SIEM, or whatever defenses an organization already has installed. Our test installation worked with Splunk, but the program also integrates with IBM's QRadar, Carbon Black, CrowdStrike, ThreatConnect, STIX and TAXII, ThreatGrid and many others, including most next generation firewalls. Once a threat is detected, BluVector can use those other programs and devices to block future, similar threats. Technically, it can and does take those future blocking actions itself, since it uses machine learning to tweak and enhance its detection engines, but it also plays well with others.

BluVector even has a Cuckoo sandbox, though it's rarely used, and only deployed when the probability engine can't get enough information from the various detection engines to make a final determination on malicious intent, a situation that we were unable to trigger. Even with its own sandbox, BluVector is designed to fully integrate with FireEye, in case an organization has already standardized on that type of sandboxing protection.

Pricing for BluVector is based on network size and whether an organization wants to fully manage the engines themselves or opt for a hosted solution where experts at BluVector help to train users, tweak the detection engines and AI, and remediate particularly thorny issues or dangerous threats found by the program. As such, pricing can vary quite a bit, but a hosted solution for a small branch office with a 500Mbps connection would run $11,400 per year, and could be considered an entry-level deployment.

Threats to networks these days are overwhelming for many organizations because of both their sophistication and volume. This is especially true for mid-tier and smaller organizations that can't invest in big cybersecurity teams, though it's a problem that affects everyone, compounded by the shortage of IT workers. BluVector can help to tame that flood, employing machines to do the heavy, cognitive work of unmasking threats at speeds that no team of humans could ever match.

When used in-line with traffic, BluVector has the potential to restore the concept of a network security perimeter, detecting previously unknown threats and stopping them before they get deep into the kill chain. Combined with machine learning that gets smarter and more network-aware over time, it can tip the scales back in favor of defenders regardless of what new malware, attack techniques or advanced threat the bad guys come up with next.

— John Breeden II

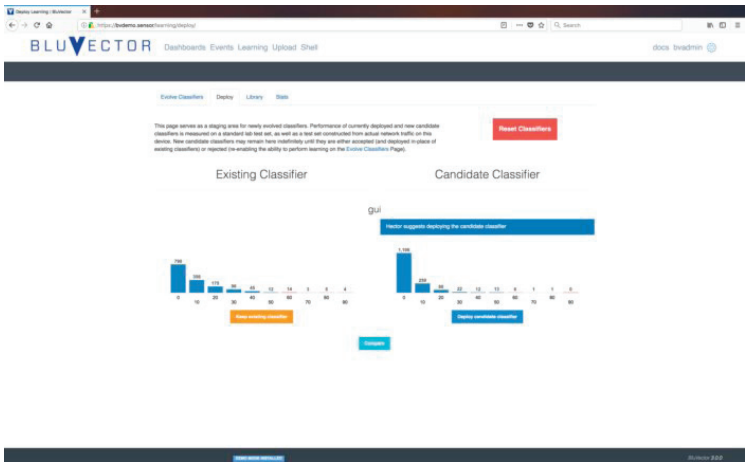

John Breeden II/IDG Users can manually tweak the way that detection engines process data, or allow BluVector to use machine learning to evolve on its own over time based on the specific network that it is tasked with protecting.

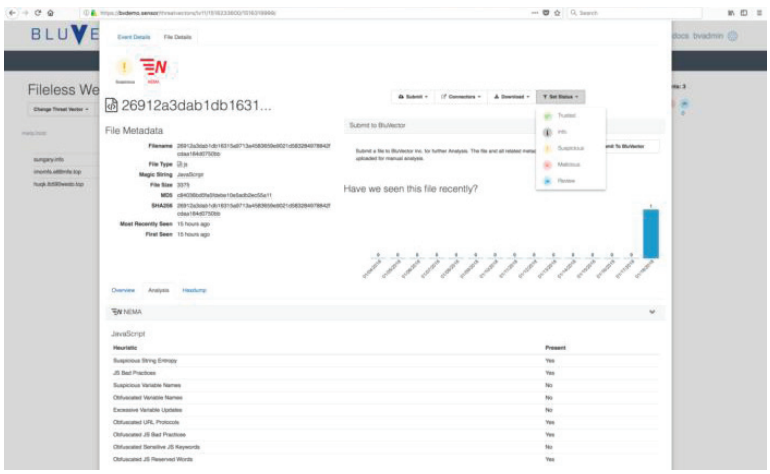

John Breeden II/IDG Even the latest malware, which is delivered as snippets of code without any associated files, can't bypass BluVector's many detection engines.

John Breeden II is an award-winning journalist and reviewer with over 20 years of experience covering technology. He is the CEO of the Tech Writers Bureau, a group that creates technological thought leadership content for organizations of all sizes.