A New Approach to Threat Detection

HOW A MASSIVE STATE MACHINE IN THE CLOUD USES REAL-TIME AND RETROSPECTIVE ANALYTICS TO IDENTIFY MULTI-STAGE ATTACKS.

A PROTECTWISE WHITEPAPER

Numerous security analytics products are trying to detect advanced threats using machine learning and big data-based architectures with large amounts of log data, metadata, and threat intel. But that's only part of the universe of relevant data. This doesn't account for clues hidden in network data from virtual environments, industrial control systems (ICS), the internet of things (IoT) and so on. Attempting threat detection without using the right kind of data yields limited or no visibility into threats.

To provide the pervasive visibility that enterprises need to combat modern attackers, a security analytics solution must handle not what is traditionally considered "Big Data", but rather a much larger set. In addition to the more traditional forms of data, a security analytics solution must be able to analyze data from extremely diverse environments, across widely distributed architectures and mesh networks, and with ephemeral services. The kinds of slowly gestating attacks being perpetrated by cyber criminals requires a security analytics solution to take a very post-big data approach, performing repeated deep packet inspection of the original raw data with the freedom to look back in time.

In addition to distilling attack insights in a post-big data world, a security analytics solution must also ensure that responders and analysts don't suffer information overload. A different approach to security visualization is needed, one that focuses on the human experience. Our world is a four dimensional one, full of complex signals highly predicated on memory. Humans have an innate ability to reason spatially, processing these complex signals to discern what is amiss. To ensure that analysts can efficiently make sense of security insights from this epic data set, a security analytics solution must deliver an immersive experience that unifies data visualization with the human ability to reason spatially. Presenting analysts with security insights in a multi-dimensional world with the ability to query across time empowers them to reason about the security realm in the same way as in the physical world.

THE PROTECTWISE GRID

ProtectWise™ provides a new utility model for security, delivering pervasive visibility, automated threat detection and unlimited forensic exploration on-demand and entirely from the cloud. The ProtectWise Grid™ enterprise security platform captures high-fidelity network traffic to create a lasting memory of enterprise network activity wherever it occurs—within the enterprise, in the cloud, in hybrid environments, or on industrial control systems (ICS).

The ProtectWise Grid ingests enterprise network activity into a finely tuned hierarchy of expert systems to analyze data and automatically detect attacks. Attacks are surfaced either by a single detection type definitively identifying a behavior or set of actions as an attack, or by a consensus between multiple, correlated detection types. The expert systems leverage a variety of technologies (e.g., machine learning, statistical modeling, heuristics, etc.) and run in parallel, without knowledge of each other. This multifaceted analysis of data means that when different techniques independently conclude that a security event has occurred, analysts can confidently assume it is something that needs to be investigated further. The breadth of the expert systems makes it possible for The ProtectWise Grid to detect both known and unknown threats.

With ProtectWise, the data stream is analyzed in real-time to identify attacks. Additionally, historical data is automatically and continually reassessed to detect prior exploits of newly discovered vulnerabilities, and integrated with the live data stream for complete context. ProtectWise is able to identify thousands of applications and protocols from the network traffic that is ingested into The ProtectWise Grid. The expert systems operate on these applications and protocols, leaving attackers nowhere to hide.

The ProtectWise Grid detects numerous categories of threats including:

- Exploits and Attacks
- Malware
- Reconnaissance
- Botnet
- Phishing
- Malicious Hosts
- APT's

All detected threat observations are automatically mapped to a specific stage of The Cyber Kill Chain. This enables security teams and incident responders to easily visualize attack progression while providing a starting point for analysis. The ProtectWise Cyber Kill Chain stages are:

- Initial reconnaissance
- Malware delivery
- Vulnerability exploit
- Victim's host beaconing
- Attacker command and control (C&C) activity
- Attacker fortification
- Attacker data theft (or other objectives including exfiltration)

The ProtectWise Grid can be thought of as a huge state machine in the cloud. This means that it can keep track of multiple sequence of events, looking back in time across all enterprise network activity. This capability is important as it enables The ProtectWise Grid to determine whether a threat observation is simply a one-off anomaly or part of a larger attack that is evolving slowly over time.

The ProtectWise Grid presents the results of analysis in a rich, innovative visualizer that delivers an immersive security experience. This gives security analysts unprecedented situational awareness and results in faster incident response and intuitive anomaly detection, even when dealing with petabytes of security data.

PROTECTWISE USES ADVANCED DETECTION TECHNOLOGIES

The ProtectWise Grid uses a hierarchy of expert systems to ensure that analysts are focusing on the threats and security events that matter to their organization. In addition to rules and heuristics, the expert systems use a variety of advanced technologies such as machine learning, intrusion detection system (IDS), automated retrospection, and customer-specific event modeling.The expert systems can be thought of as nodes in an artificial neural network (ANN), all working in concert to identity threats, with low false positive rates.

The expert systems attribute a threat score, ranging from 0 to 100, to the most severe part of the activity associated with an observation or security event. The severity of an observation or security event is directly proportional to the threat score and can be used by analysts to prioritize their investigative efforts. The diagram below provides a view into some of the expert systems and the range of scores associated with the triggering of those systems.

Additionally, new approaches to threat detection are continually being integrated into The ProtectWise Grid, enabling it to stay ahead as attacker's techniques mature.

MACHINE LEARNING

The ProtectWise Grid uses a number of machine learning algorithms, not just for attack detection but also for classification and data enrichment. Many of the machine learning models are trained and updated across the breadth of data under ProtectWise's purview and are not necessarily specific to a given customer's individual data, meaning that even new customers can take advantage of ProtectWise's machine learning capabilities right out of the box. Models are thereafter customized and further refined within a specific customer's data silo as data comes online, but each model can start from an informed first estimate using the scale of ProtectWise's monitoring experiences.

Hierarchy of Expert Systems

The ProtectWise Grid uses a variety of expert systems for threat detection, leveraging advanced technologies (e.g., machine learning, statistical modeling, intrusion detection, correlation, rules and heuristics) to produce credible results. The ProtectWise Grid takes a hierarchical approach with its expert systems when elevating a threat observation or series of observations to the status of a security event, which is where security analyst's investigative efforts should be focused.

With ProtectWise, a security event is generated either by a single expert system definitively identifying an action or set of actions as an attack (e.g., finding malware based on an IDS signature), or by a consensus between multiple expert systems, each using a different technique for independent analysis. For example, the IDS identifying malware based on a signature match is all that's needed to generate a security event for a known threat. But when dealing with unknown threats (e.g., a multi-stage attack) a consensus between multiple expert systems is required before ProtectWise generates a security event.

The hierarchy of expert systems provide security analysts with the assurance that security events generated by The ProtectWise Grid are indeed real, and not simply alarms about suspicious activities that add to alert white noise and contribute to analyst fatigue.

Exploit

1 event from 20:49:20 to 20:49:28
on 2017-01-10

Delivery

1 event from 20:49:20 to 20:49:28
on 2017-01-10

Beacon

3 events from 20:49:20 to 20:49:28
on 2017-01-10

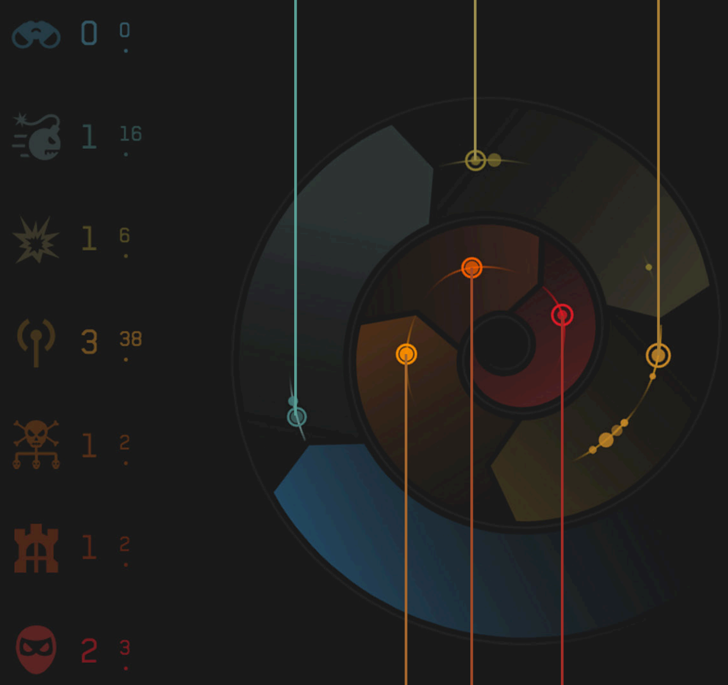

ProtectWise uses the Attack Spiral to display attack progression by The Cyber Kill Chain stages, providing a quick visual summary of the variety of threats and a sense of severity. The Attack Spiral is easy to interpret. Threat observations or events are located on the outer edge of the attack spiral (e.g., blue or yellow which represents reconnaissance or beaconing) are not as severe as those near the center (e.g., red represents data theft).

C&C

1 event from 20:52:16 to 20:52:24
on 2017-01-10

Data Theft

2 events from 20:49:44 to 20:49:52
on 2017-01-10

Fortification

1 event from 20:51:12 to 20:51:20
on 2017-01-10

Supervised Machine Learning Algorithms.

ProtectWise data scientists train these algorithms to differentiate between security events and non-events by providing labeled training data. For example, some malware uses DGA (domain generation algorithms) to generate many domain names which are used as communication points with C&C (command and control) servers. ProtectWise's data scientists curate large amounts of incoming data to find statistically high-quality training data, enabling the DGA module to learn the difference between a "real" domain (e.g., cnn.com) and a "generated" domain (e.g., wyzngmatzmg.com). Once trained, the DGA module is used to identify algorithmically generated domains that arrive as part of new incoming netflows, and attaches that information as context to its event engine, where it can be used to determine the existence of possible threats to the network. ProtectWise currently uses supervised machine learning algorithms to identify malicious activities such as DGA, DNS tunneling and malicious certificates, and has plans to implement these techniques across a larger array of security problems in the future.

Unsupervised Machine Learning Algorithms.

The ProtectWise Grid makes extensive use of unsupervised machine learning techniques, which is used to provide functionality that produces results visible to analysts (e.g., anomaly detection) and for behind-the-scenes activities that help inform other parts of the platform.

For example, in anomaly detection, The ProtectWise Grid learns what constitutes "normal behavior" relative to specific markers or indicators within each customer's environment (and possibly for each device within the environment). The algorithms then detect anomalies that are indicative of attacks based on statistically significant deviations from the derived behavioral patterns. ProtectWise will also continually rebaseline behaviors as "normal behavior" may change over time.

The behind-the-scenes activities in which ProtectWise employs unsupervised machine learning include tasks such as asset classification/grouping and information concentration.

- For asset classification tasks, The ProtectWise Grid sub-groups assets on the network based on their behaviors and learns what to expect to from those assets groups. For example, the behaviors (e.g., the software used, the sites visited, etc.) seen on machines used by marketing personnel would be different from those used by executives, which would be different than those used by engineers, and so on. Unsupervised machine learning equips The ProtectWise Grid to cluster machines based on patterns of historical behaviors and create more appropriate behavioral baselines for each user-group, rather than simply employing a single baseline in attempting to describe a large number of possibly widely-varying behaviors.

- The ProtectWise Grid also uses unsupervised techniques for information concentration, also known as dimensionality reduction. This can be divided into feature selection, which is the process of selecting the variables used in constructing a machine learning model, and feature extraction, which is the process of starting with observed data and building derived values that facilitate learning and generalizing of the models. Dimensionality reduction increases the efficiency of machine learning models and can also be used to discover latent patterns to inform others models.

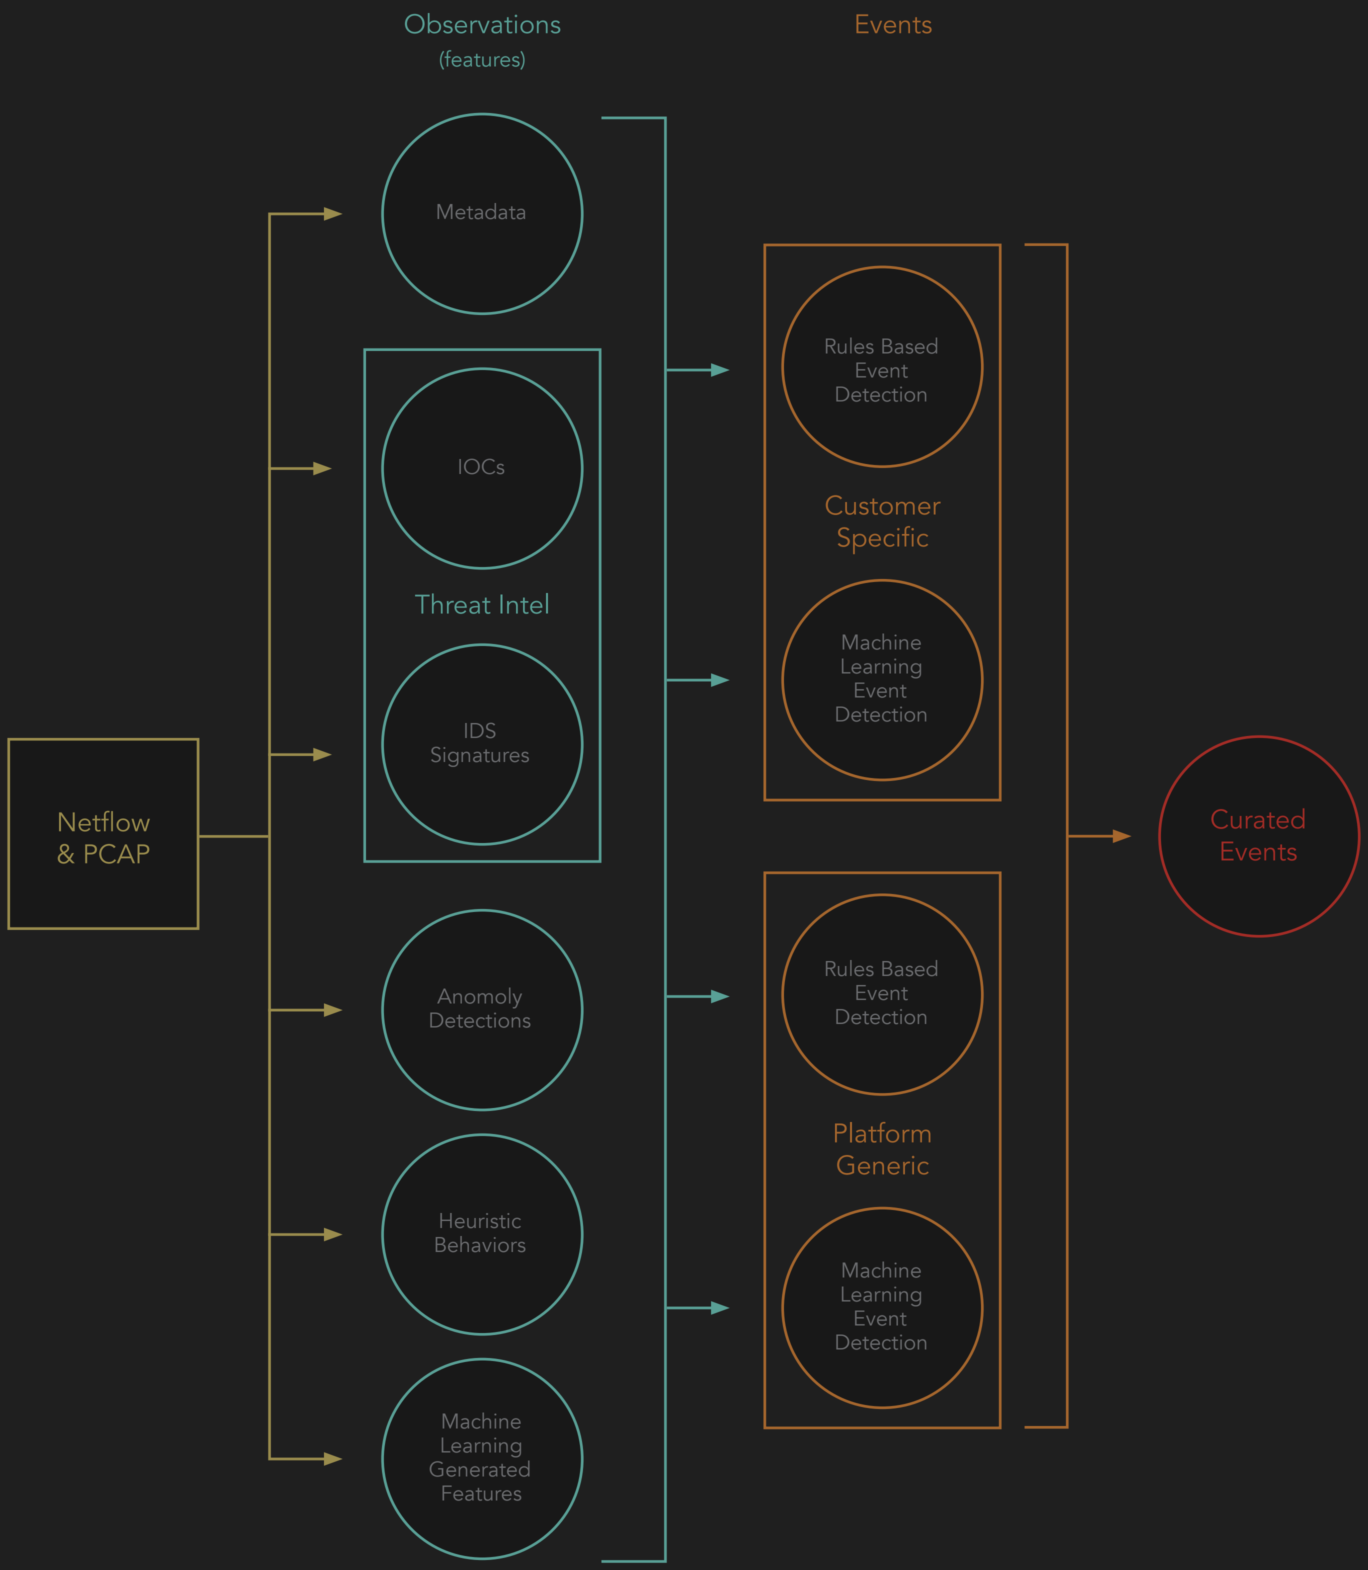

The expert systems in the The ProtectWise Grid as can be thought of as nodes in an artificial neural network (ANN).

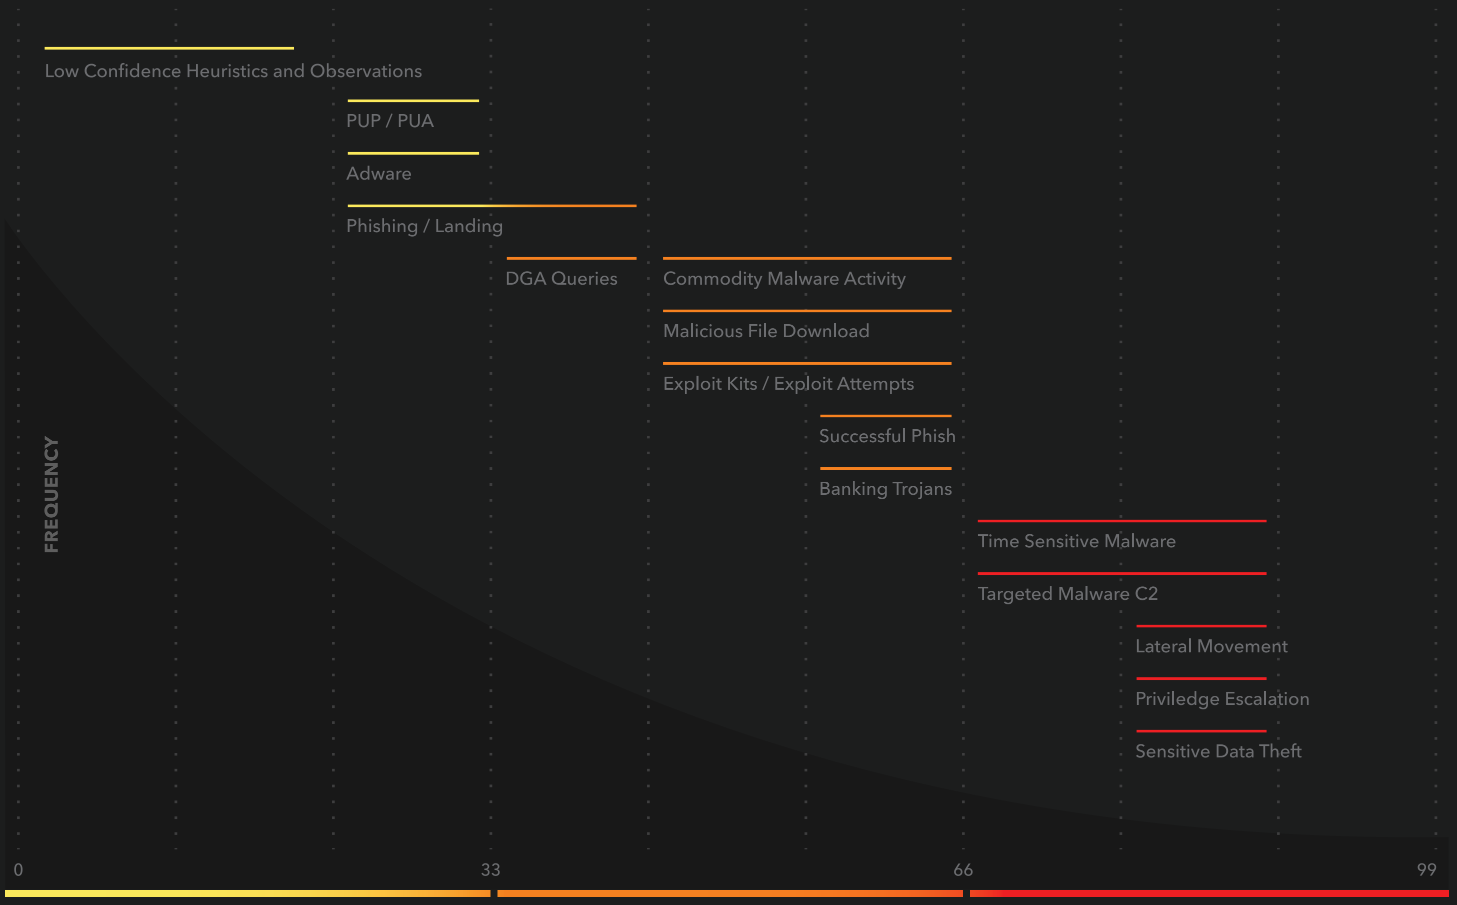

LOW
Low risk, low confidence, possibly malicious, possible policy violations.

MEDIUM
Risk likely isolated to a single host or user, victims are likely to be targets of opportunity.

HIGH
Risk is to the organization, activity is likely targeted.

Some of the expert systems used by The ProtectWise Grid and the range of associated threat scores that are attributed to threat observations and security events generated by these expert systems.

INTRUSION DETECTION (IDS)

The ProtectWise Grid includes a full intrusion detection system (IDS) engine. Using IDS as one of its expert systems enables The ProtectWise Grid to deterministically detect the types of threats (e.g., malware) that can be identified via signatures and elevate those to security events.

The ProtectWise Grid includes more than a dozen carefully selected and curated threat intelligence sources (open source, fee-based) and reputation lists (IP, DNS, URL, File, and Certificate) which are constantly updated and maintained by ProtectWise. This ensures the cleanest and most up to date threat intelligence and accurate detection with the IDS. Customers can also use threat intel from external sources (e.g., Facebook threat exchange, etc.) and in-house research in the IDS.

The ProtectWise Grid is also able to accommodate any number of customer-supplied custom threat intel sources. The intent of 'Bring Your Own Intelligence' (or BYOI) is to allow customers to modify and add their own threat intelligence from many other sources, including law enforcement, industry consortiums, internal research/threat hunting or 3rd party intelligence feeds. The ProtectWise Grid can use this intelligence for real time as well as retrospective detection. With BYOI, customers are not restricted to a one-size fits all approach which can be ineffective for those customers with diverse networks. By accommodating BYOI, The ProtectWise Grid enables more effective threat detection as customers can leverage the threat intelligence they have that's uniquely available to their organization while suppressing other intelligence that is ineffective in their network.

The included IDS means that enterprises can replace their ageing IDS appliances with The ProtectWise Grid and even import current IDS rules. Plus with The ProtectWise Grid, enterprises get a scalable cloud platform with a richer set of other security capabilities such as network traffic recording, retrospective analysis, multi-stage attack detection and more.

Curation of Intel

The ProtectWise Threat Research and Analysis team curates and monitors security events detected across ProtectWise's customer base, while continually refining and enriching applied threat intelligence and complementing it with proprietary heuristics and classifiers.

ProtectWise applies threat intelligence from a variety of sources including third party vendors, open source intelligence, both closed and open sharing groups, and in-house security research. In the case of indicators that are low in volume and specific in context, such as IDS rules, the ProtectWise threat research reviews and tests them against existing traffic for a period of time before making them available to customers in The ProtectWise Grid. The ProtectWise Threat Research and Analysis team continually monitors the performance of external threat intel sources and maintains a whitelist of reputable sources. With open sharing groups, indicators are more critically evaluated until ProtectWise can determine the trustworthiness of the source.

In addition to intel feeds that are curated and monitored, the ProtectWise Threat Research and Analysis team derives indicators based on classifiers and heuristics, such as Machine Generated Domains and SSL certificates, or examination of a domains age. These are pre-vetted with a whitelist and then monitored for performance. The end result is a refined detection signal that correlates multiple types of intelligence and indicators. The ProtectWise Threat Research and Analysis team continuously refines this signal looking across all security events, in real time as well as retrospectively.

RETROSPECTIVE ANALYSIS

The ProtectWise Grid also provides retrospective analysis. which is the ability to go back in time (by looking at historical data) to uncover evidence of previously unknown threats as new attack campaigns are discovered. With ProtectWise, retrospection can be manual or automatic. Retrospection is only possible because The ProtectWise Grid creates a perfect memory of your enterprise network traffic, retaining full PCAP data.

During automated retrospection, The ProtectWise Grid conducts regular scans of historical data in customer environments to see if there are any matches with newly available threat intelligence and generates security events as necessary. If new campaigns are discovered in historical data, the full PCAP makes it easy for analysts to validate attacks, identify victims and reconstruct the resources used in the attack.

CUSTOMER-SPECIFIC EVENT MODELING

What may be a security event for one customer could be an insignificant anomaly for another. Recognizing this difference, The ProtectWise Grid does not take a “one size fits all” approach when identifying security events. Rather, a combination of fully supervised and unsupervised machine learning techniques, both firing at the same time, are used to elevate threat observations to security events, based on a customer’s needs. By learning what constitutes higher value, lower noise security events on a per customer basis, ProtectWise enables enterprises to get better productivity out of their security teams.

Supervised Event Modeling.
ProtectWise threat researchers regularly examine how customers are using The ProtectWise Grid and periodically generates an unique event model for each customer environment. These machine learning models are published to production in the form similar to a TTP (tactics, techniques and procedures) library and determine the variety of threat observations that will get elevated as security events for a particular customer.

Unsupervised Event Modeling.
The ProtectWise Grid event engine uses unsupervised machine learning to automatically make similar kinds of determinations about what constitutes an event in a customer’s environment. This happens completely autonomously and is made possible because The ProtectWise Grid learns from how customers use data, continually observing the labeling that’s happening.

HEURISTICS

On the spectrum of threat detection techniques, heuristics falls in the middle between traditional IDS / threat intelligence and machine learning / behavioral analysis. Heuristics are similar to traditional IDS signatures in that they are deterministic rules defined by a security analyst. However they differ from traditional IDS signatures as they are looking at more loosely defined behaviors. The ProtectWise heuristic engine keeps state over a long period of time, allowing the heuristics to watch for different behaviors unfolding across a large amount of data points. ProtectWise uses heuristics to create threat observations for activity that is difficult to detect using traditional IDS signatures, in particular those that do not have unique content in the packet payload, such as port scanning, brute force password attempts, protocol abuse, and so on. ProtectWise also uses heuristics to lift and correlate observations that are part of a larger security event. An example of this is a heuristic that watches for successful exploitation that ultimately progresses along the Cyber Kill Chain, such as a malicious file download followed by beaconing and ultimately command and control several hours later.

CONCLUSION

The variety of approaches to threat detection (e.g., machine learning, IDS, threat research) equips The ProtectWise Grid to discover both known and unknown attacks. Requiring consensus from the expert systems prior to generating security events reduces the false positives. A broad coverage model enables The ProtectWise Grid to provide visibility into attacks wherever they are happening—within the enterprise, in the cloud, in hybrid environments, or on industrial control systems (ICS). The combination of a broad coverage model, multi-faceted analysis and consensus-based generation of security events means that The ProtectWise Grid leaves attackers with nowhere to hide and security analysts can be more productive.

Analysis is performed in real-time on live data, with historical data brought in for context and continuous reassessment. Near real-time detection reduces attack dwell time (i.e., time from infection to detection), and limits the extent of the damage that can be inflicted. Other solutions, such as those architected using big data, can't detect in near real time due to the batch nature of their analysis.

The ProtectWise Grid takes advantage of its cloud form factor to take the learnings from one enterprise network and create a shared wisdom across all. The cloud also lets The ProtectWise Grid maintain a perfect memory of the enterprise network, potentially for an unlimited amount of time. This allows enterprises to more confidently (more than any other in-market solution) determine whether or not they've been ever been affected when news of a new zero day attack breaks or if they discover some new threat as their security team threat hunts.

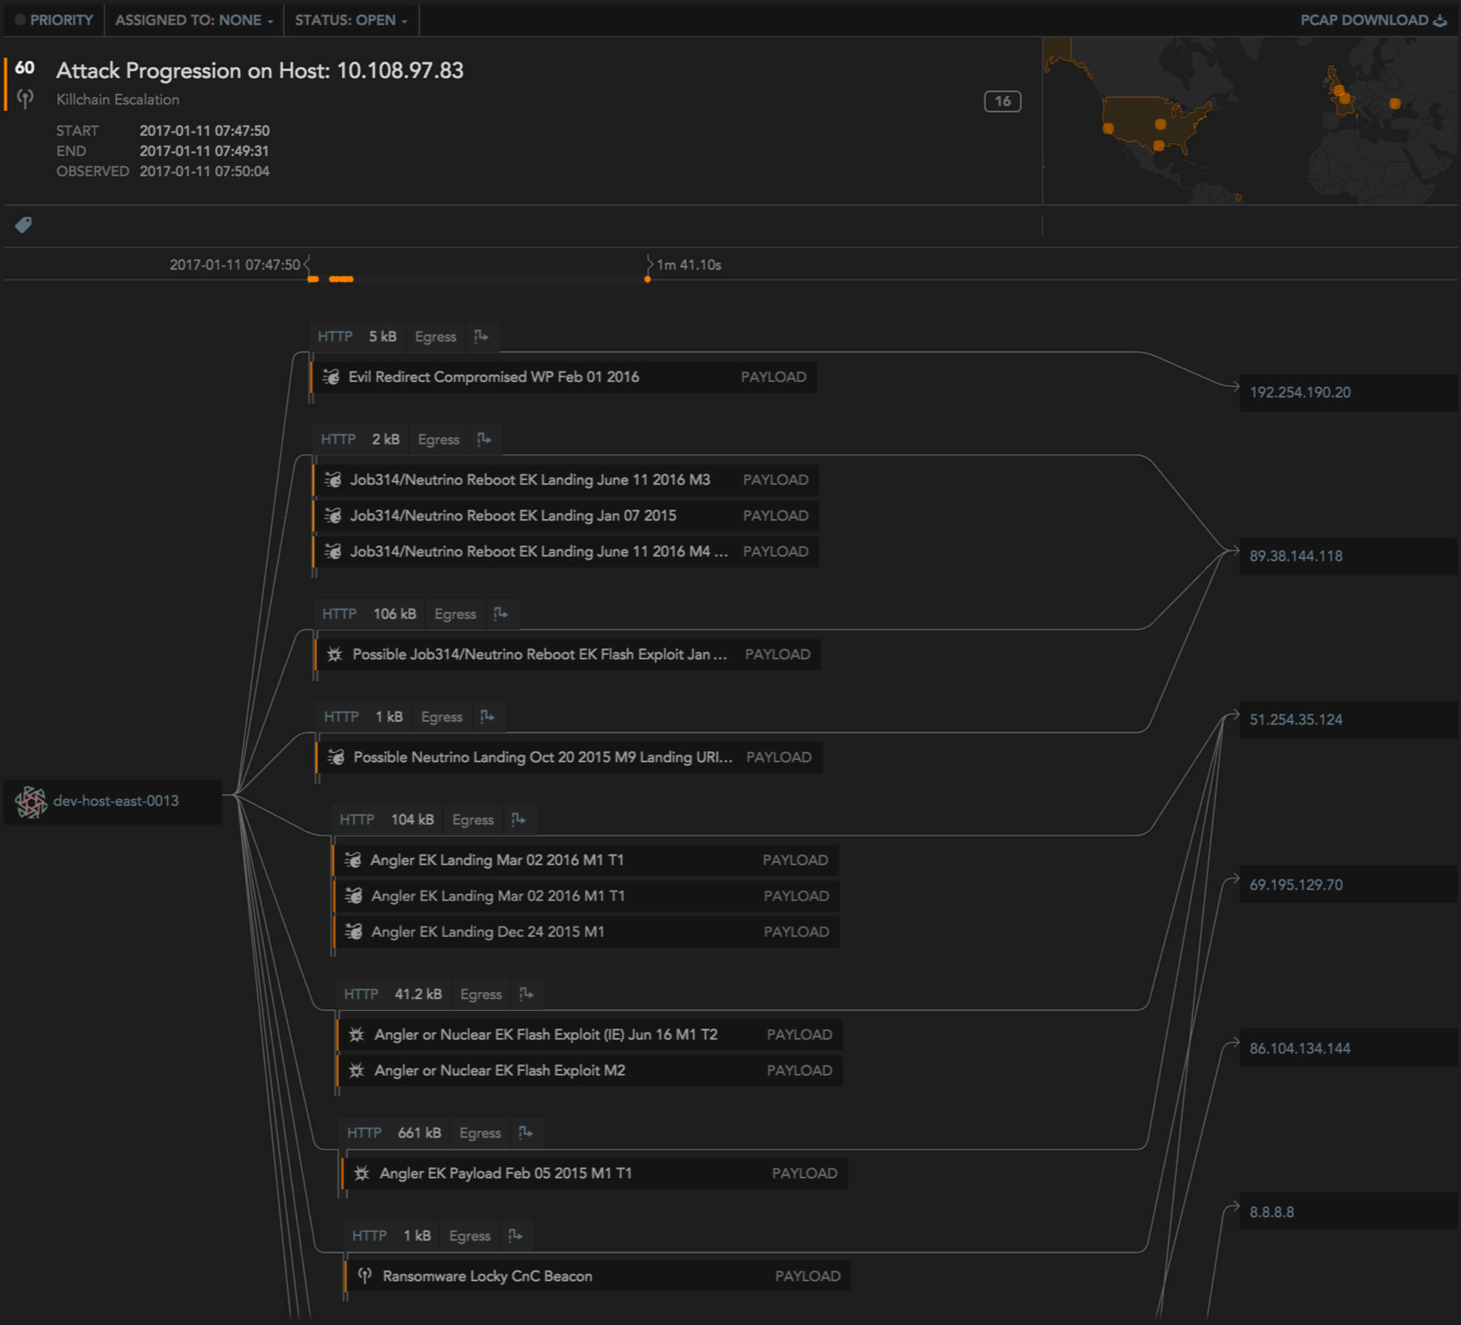

PRIORITY | ASSIGNED TO: NONE | STATUS: OPEN | PCAP DOWNLOAD

60

Attack Progression on Host: 10.108.97.83

Killchain Escalation

START 2017-01-11 07:47:50
END 2017-01-11 07:49:31
OBSERVED 2017-01-11 07:50:04

16

2017-01-11 07:47:50 — 1m 41.10s

dev-host-east-0013

HTTP 5 kB Egress
- Evil Redirect Compromised WP Feb 01 2016 — PAYLOAD

HTTP 2 kB Egress
- Job314/Neutrino Reboot EK Landing June 11 2016 M3 — PAYLOAD
- Job314/Neutrino Reboot EK Landing Jan 07 2015 — PAYLOAD
- Job314/Neutrino Reboot EK Landing June 11 2016 M4 ... — PAYLOAD

HTTP 106 kB Egress
- Possible Job314/Neutrino Reboot EK Flash Exploit Jan ... — PAYLOAD

HTTP 1 kB Egress
- Possible Neutrino Landing Oct 20 2015 M9 Landing URI... — PAYLOAD

HTTP 104 kB Egress
- Angler EK Landing Mar 02 2016 M1 T1 — PAYLOAD
- Angler EK Landing Mar 02 2016 M1 T1 — PAYLOAD
- Angler EK Landing Dec 24 2015 M1 — PAYLOAD

HTTP 41.2 kB Egress
- Angler or Nuclear EK Flash Exploit (IE) Jun 16 M1 T2 — PAYLOAD
- Angler or Nuclear EK Flash Exploit M2 — PAYLOAD

HTTP 661 kB Egress
- Angler EK Payload Feb 05 2015 M1 T1 — PAYLOAD

HTTP 1 kB Egress
- Ransomware Locky CnC Beacon — PAYLOAD

192.254.190.20
89.38.144.118
51.254.35.124
69.195.129.70
86.104.134.144
8.8.8.8

The multifaceted analysis performed by expert systems enable The ProtectWise Grid to take seemingly disparate threat observations and correlate them to relevant security events, providing security analysts with complete visibility into what triggered the event.

About ProtectWise

ProtectWise™ is disrupting the security industry with The ProtectWise Grid™, its enterprise security platform that captures high fidelity network traffic, creates a lasting memory for the network, and delivers real time and retrospective alerting and analysis in a rich, innovative visualizer. By harnessing the power of the cloud, The ProtectWise Grid provides an integrated solution with complete detection and visibility of enterprise threats and accelerated incident response. The ProtectWise Grid delivers unique advantages over current network security solutions, including an unlimited retention window with full-fidelity forensic capacity, the industry's only automated smart retrospection, advanced security visualization, and the ease and cost-savings of an on-demand deployment model. For more information, visit www.protectwise.com.

20170117

PROTECTWISE™

1601 WEWATTA ST • SUITE 700 • DENVER, CO 80202 • 1.303.625.6802
WWW.PROTECTWISE.COM • INFO@PROTECTWISE.COM