SIGNALSENSE

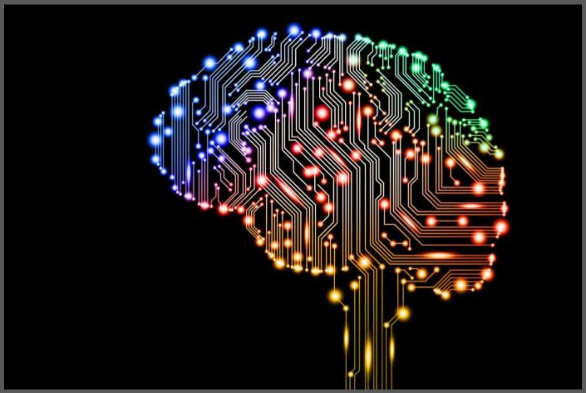

White Paper

USING DEEP LEARNING TO DETECT THREATS

Executive Summary

Everyday the headlines reflect the severity, and challenge of data breaches and network compromises from APT's or Advanced Persistent Threats. Nobody is immune. Inside risks, threats and vulnerabilities exist today within most enterprises, and vigilant defense and adaptive detection requires a new approach and paradigm.

Traditional security monitoring solutions rely on brittle, signature based perimeter defenses, which today's sophisticated malware can evade through a variety of methods. Furthermore, in an effort to enable collaboration and agility across today's enterprise, the perimeter is dissolving quickly. Also, threats and risks continually change and transform, so keeping perimeter checkpoints up to date can be an arduous task for any service provider or organization.

On their own, low-level behaviors, or signals, and subtle indicators would go unnoticed, and are difficult to detect and investigate at scale cost effectively. Yet in aggregate, these can signify a salient risk, Indication of Threat or Compromise.

The ramifications associated with a breach can be enduring; reputation, top line revenues and many other undesirable outcomes ensue. Most organizations realize their network is compromised, and instead are focusing on minimizing the consequences associated with an inevitable breach.

Addressing these threats requires a new approach and paradigm for network security, continuously monitoring, adapting and learning through experience, improving detection accuracy, reducing "false positives", and arming your security team with real time prescriptive guidance on which priorities to investigate is essential.

Evidence? Consider this…

Whether from third party actors, policy negligence, or malicious intent, threats exist inside the network. In fact, most security executives acknowledge their network has already most likely been compromised, which means the key is early proactive and adaptive detection, coupled with mitigating the consequences associated with a breach.

One key report, the Global State of Information Security Survey 2015 from PwC, found that the highest number of perpetrators of insider crimes were current employees (32%) followed by former employees (30%). Other perpetrators include partners, contractors and customers. While the 2015 Data Breach Investigations Report from Verizon adds that of insider incidents that occur, 37.6% were from ordinary end users, 16.8% were from cashiers, 11.2% from finance staff and 10.4% from executives.

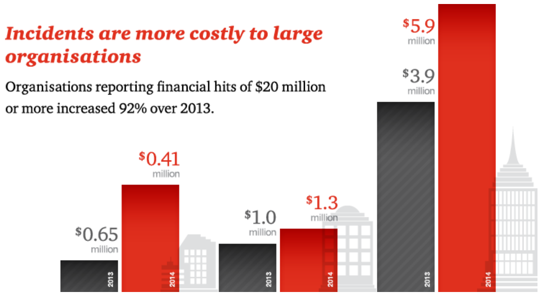
Incidents are more costly to large organisations

Organisations reporting financial hits of $20 million or more increased 92% over 2013.

If budgets represent priorities, we've definitely got a problem: The PwC survey comments that crimes caused by internal perpetrators are often more costly or damaging than those perpetrated by external groups. The SANS report adds that while more than half (52%) of respondents perceive negligent employees as the cause of significant damage, almost half (44%) are spending 10% or less of their IT budget on this insider threats, “so it’s clear why survey respondents also suffer a significant number of insider breaches.”

Get “Clarity” about your Network Traffic

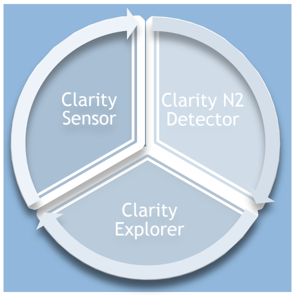

SignalSense was created to address the aforementioned challenges; the growing need to continuously monitor and visualize network traffic from the inside out, leveraging Deep Learning to create a “Neural Network” trained on a specific environment that can adapt, refine and control the accuracy of machine-identified threats, saving valuable time to detection and remediation.

This requires the combination of several solutions, including a scalable sensor to capture and record all network traffic, providing intelligent inline classifications of endpoints, host names, IP reputations, packet and flow data and other key attributes, coupled with a capability to use Deep Learning, what we call “N2 Detector”, to baseline these environmental findings, calibration to identify deviations from what are learned as “normal” behaviors for a “typical” environment, and, finally, visualization of all of this, so a security professional can immediately pivot into investigation of prioritized anomalies, IOC’s and IOT’s.

Our suite is called **“Clarity”**; which captures, identifies and adapts to your unique network environment. We are the first to combine all three key products into a single solution. **Clarity: Capture, Detect and Explore**.

Check your “Blind Spot”

As mentioned earlier, today’s standard security approach to perimeter detection is to monitor and keep known threats outside from entering your network, using next generation Firewalls that incorporate IPS/IDS, WAF and other more intelligent means of identifying patterns and new signatures at the perimeter.

Unfortunately, as we have discussed, many of today’s breaches are caused from inside actors, opening up a port for communication, or simply connecting a mobile phone, USB or wireless device to the network.

Modern Malware and Threats can evade traditional signature based methods by obfuscating or encoding the payload. Encrypted protocols and polymorphic code can circumvent or penetrate today’s next generation firewall defenses as well.

In addition, it is very difficult to establish what “normal” behavior represents without a Deep Learning engine to help characterize these anomaly types and classes. Most Endpoint, IPS/IDS and other Security Analytics tools send thousands of alerts to Security Incident and Event Management solutions (SIEM), which make the efficacy of these dashboards challenging, given the time it takes to eliminate benign alerts (or signals) and prioritize the real threats to investigate.

The Age of Machine Intelligence

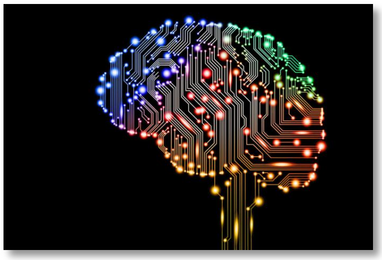

Security approaches that rely on manually filtering and investigating alerts don't scale well: high numbers of false positives consume ever-increasing time, and there's always a shortage of expertise to investigate incidents. Expanding the horizons of security threat detection to find advanced attacks generates vastly more data and more complexity. Attacks that involve multiple tactics, tools, and endpoints throughout an enterprise demand a new level of detection. Attack and anomaly detections must consider hundreds of indicators and interaction patterns across thousands of machines on a continuous basis. Further, because attackers are constantly changing their techniques faster than manually driven forensics and threat intelligence can keep up; detection must be adaptive and generalize with "intuitive knowledge" just as the best security analysts do.

SignalSense's vision is to expand the scope and depth of security detection using machine intelligence. Deep Learning neural nets have a remarkable capacity for teasing out and making the patterns hidden within terabytes of data, and Deep Learning nets excel at problems that require human-level complex perception like computer vision, speech recognition, and natural language processing. SignalSense Clarity N2 Detector uses Deep Learning to take low-level data and precisely identify protocols, endpoint behaviors, and known threats, all of which are critical context for security analysts.

Clarity N2 Detector then takes this context at a macro level to consider the millions of interactions between endpoints, learning how a company's systems and employees interact. With this knowledge, Clarity N2 Detector uncovers anomalous behavior patterns and activity that indicates complex, remotely directed external attacks as well as insider threats. And since what is anomalous vs. what is acceptable behavior is a complex judgment, one that is constantly changing as business & attackers evolve, Clarity N2 Detector is designed to learn and improve on an ongoing basis.

Using the Clarity Sensor's high performance GPU computing capacity, Clarity N2 Detector continuously re-analyzes your company's systems and interactions and updates detection models, keeping accuracy high and avoiding expensive "baselining" requirements. Further, Clarity N2 Detector leverages feedback from security analyst users to learn the subtle judgments that are the difference between potential threat and false positive. By learning by example, SignalSense tunes anomaly detection and threat detection to understand a company's unique requirements and operations. Security analysts get high quality threat detection based on a much broader horizon, so that they can focus on previously unknown attacks and assess risks much more quickly.

Anomaly Examples; The Case for Deep Learning

Compromised VOIP Control System

SignalSense Clarity Suite identified a compromised on-premise VOIP control system, which was managed by a vendor. SignalSense's neural net-based detectors identify the application protocols in use in IP flows, even flows that are encrypted such as SSH. The customer company's VOIP system was showing an unusually high number of successful SSH connections, and further those connections were with external IP addresses in China. SignalSense Clarity Suite identified the encrypted connections, the high volume of flows, flagged external IP addresses with reputation alerts, and identified the suspect geography. The anomalous activity was clearly presented using Clarity Explorer's visualization of anomalies and quickly localized to a compromised OS installation on the VOIP system.

After the VOIP management vendor was notified and they re-imaged the system, within minutes SignalSense again detected that the VOIP server was still compromised, indicating that the vendor's authorization credentials and/or OS image were compromised prior to installation. The customer escalated the issue to the device OEM and used Clarity Suite's investigative console to download hundreds of example SSH flows as pcap files to support the OEM's internal investigation.

SignalSense identified internal network scanning, where one endpoint was looking for exposed services with known vulnerabilities and performing dictionary attacks on remote access services. SignalSense's N2 Detection Engine flagged unexpectedly high numbers of destination endpoints and destination ports accessed by the scanning machine, outside the normal profile for a workstation machine on the company network. The traffic was also flagged as unusual based on the scanning machine's own prior history over the past month. Network scanning is an indicator of compromise by an attacker exploring and mapping a company's resources, and also an indicator of potential insider threat.

Previously Unseen Application Protocols

SignalSense's neural net-based detection engine learns the expected network protocols and sending applications that appear in a company's traffic mix. This learning, or "baselining" process takes place over several days, where SignalSense trains a detection model on the protocols in use and generalizes observed variations on those protocols. Once trained, SignalSense's N2 Detection Engine models identify previously unseen protocols and protocol usage that does not fall within the expected patterns. When new applications appear on the network, due to unmanaged mobile device apps or unmanaged applications downloaded by end users onto workstations, SignalSense flags novel protocols and usage of protocols and identifies the devices sending those protocols. Using Clarity Explorer, the device behavior is examined and protocols dissected to determine whether there is a threat.

SignalSense continually updates its trained models to account for changes in a company's compute environment, and learns from security analyst feedback on false positives so as to not repeatedly identify novel but benign protocols.

Unexpected Structural Anomalies in Packets and Flows

SignalSense N2 Detectors learn a company's network behaviors at several levels – behavior between endpoints, with external services, and at a structural level. When attackers take advantage of software vulnerabilities, such as the recently OpenSSL Heartbleed Vulnerability, those attacks often involves creating packets or flows that are altered to trigger the vulnerability, yet the packets are still valid on the network. Deep Packet Inspection often passes such packets.

SignalSense N2 Detectors find novel, altered packets that don't match prior learned examples from the company's network. A pattern of repeated altered packets may indicate malfunctioning equipment or the emergence of a new type of attack.

Potential Insider Threat

A customer endpoint began making long, repeated file transfers that were out of character with the prior history of the endpoint. SignalSense flagged this behavior as an anomaly because of the length and quantity of the file transfer data flows. The N2 Detection engine additionally detected the protocol of the flows and behavior class of the endpoint, providing fast context during the investigation process. Such activity can indicate insider threats or the presence of undetected malware exfiltrating data.

Applications in Violation of Policy - BitTorrent and Tor

SignalSense identified customer endpoints running software that was in violation of policy. SignalSense's N2 Detection engine detected the presence of encrypted data and identified the streams as being sent by the Tor Browser application. Use of Tor was against corporate policy, and further has been observed as a command and control transport mechanism by some families of malware.

SignalSense also identified endpoints running BitTorrent nodes, also in violation of company policy. The network behavior of connection attempts to external endpoints that did not respond triggered anomalous behavior events from SignalSense's anomaly detection models. The anomalous behavior was confirmed by using Clarity Explorer visualizations showing the geographic diversity of the output connection requests, and secondarily for the connections that were successful, SignalSense N2 Detection models for predicted protocol identified the protocol as BitTorrent traffic.

Improperly Configured Switch Rules Allowing Inbound Internet Scans to Workstations

A SignalSense customer deployed a SignalSense sensor. Within an hour of data gathering, Clarity Suite identified unexpected internet traffic inbound directly to workstations that sat behind the company's firewall. Clarity Suite's visualization tools include pre-built common reports to find inbound traffic, and using Clarity's machine-learning based classification of device behavioral profiles it was easy to segment the inbound traffic into a bucket for expected traffic to internet-exposed servers vs. unexpected traffic to workstations. Further the inbound traffic was flagged as anomalous because it was using the SSH protocol, which is not an expected protocol for Windows workstations. The customer was able to identify the subset of internal endpoints exposed to scanning traffic, and traced the root cause to a misconfigured switch.

New Variations on Known Threats

New versions of existing malware families are continuously being created and propagated. IOCs based on file signatures, Windows registry entries, and processes become out of date quickly as malware authors tweak malware executables to evade detection. SignalSense N2 known threat detection models identify malware using learned, generalized network behavior patterns rather than signatures that may be evaded if one byte gets changed. By contrast malware command and control protocols require more significant investment by malware authors to change, and therefore evolve more slowly. SignalSense N2 Models identified CNC flows and behaviors even though the delivery and packaging of the malware changed to evade signature based detection.

Conclusion

In an age of ubiquitous threats and attacks, enterprises must become more vigilant about monitoring traffic both inside and outside of their networks. Learning through experience can be applied to help fortify the adaptive and proactive detection, prescriptive response and overall security fabric of a company, the same way that this breakthrough technology is used today for facial recognition. Deep Learning is the next wave of computing, and age of Machine Intelligence.

Accepting the reality that risks and threats are persistent and exist today within your enterprise, and that regardless of how well informed and comprehensive your security policy is, compromise is inevitable, then leveraging solutions that can monitor, adapt and visualize these events, in real time, including behaviors, patterns and flows are essential to early detection and remediation.

By combining the capture, enrichment, learning and visualization of anomalies, patterns and behaviors across all network traffic, SignalSense can quickly identify the right indicators of threats or compromise in real time, while preserving valuable investigation time and resources within your security organization.

SignalSense: Visibility. Insight. Security.