Network Behavioral Analytics with the Securli® Advanced Threat Defense Platform

Contextual Big Data to Counter Attack-in-Depth Invasions

Big data is the current rage in cyber security. There is a general perception that throwing lots of network data at a big data engine is the way to identify malicious behaviors. There are two significant problems with this theory:

1. A big data analytics tool is only as good as the content from the data sources that feed it, and
2. Analysis without context fails to establish threat relevance and is useless for defense, detection and remediation.

Typical data sources such as log files, NetFlow and baselines are missing all of the key indicators of malicious behaviors, and instead depict activity that appears to the typical data analytics engines as seemingly benign traffic. As malware continues to evolve and insiders are now operating largely in stealth mode and understanding the constructs of these data analytics engines, fewer and fewer of the recognized indicative data elements are showing up in these logs, flows and baselines.

In addition, today's coordinated attacks are multi-stage and multi-vector. But because traditional big data analytics examines discrete events out of context they miss the subtle patterns and sequences of related behaviors that cyber-criminals are now using consistently across the global threat landscape to assemble an effective Attack-in-Depth invasion model. Attack-in-Depth is a summary version of the once popular cyber kill-chain model that works by delivering payloads, persisting on endpoints, taking hold across the network and exfiltrating or destroying information assets.

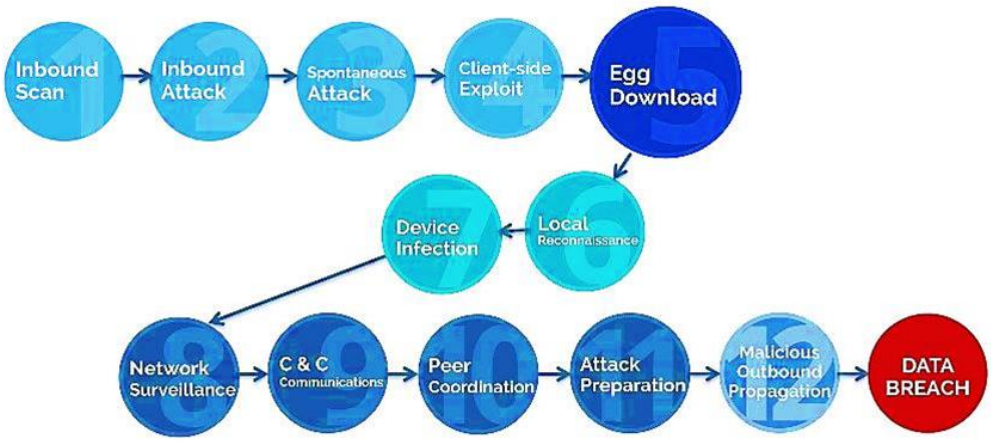

To successfully combat these Attack-in-Depth threats we must shift our approach to contextual data analytics.

An effective analytic engine must be fed the otherwise hidden indicators of malicious behaviors, indicators that are only detected with the right type of analytics.

These analytic engines must use designed algorithms that are constructed to detect both structured and unstructured malicious behaviors within the context of a specific threat envelope. That threat envelope must be informed by patterns of behavior occurring outside the network and across a spectrum of threat landscape external to the operation. And, these engines must be able to operate on this data in real time to identify and isolate an infection after a network has been invaded and before the assets can be breached. This can only happen with contextual analytics.

A Review of Different Analytic Approaches

At their core, analytics engines typically follow one of four primary reasoning methodologies:

Deductive Reasoning – Deductive reasoning is based in the theory of deductive inference that draws specific conclusions from general rules e.g., If A = B and B = C, then A = C, regardless of what A or B contain. Deductive reasoning tracks from a general rule to a specific conclusion. If original assertions are true then the conclusion must be true. A fundamental weakness of deductive reasoning is it's often Tautological (e.g. Malware contains malicious code and is always true) and it is unaffected by contextual inputs, e.g., to earn a master's degree, a student must have 32 credits. Tim has 40 credits, so Tim will earn a master's degree, except when he decides not to.

In security analytics, A only equals B most of the time and sometimes it can equal D, so A cannot always equal C, therefore using deductive reasoning as a basis for detection analytics is a flawed way to try and predict the future. You are theoretically guaranteed to be wrong at least once.

In general, common signature-based systems such as IDS/IPS and endpoint security are deductive in nature.

Inductive Reasoning – Inductive reasoning is the opposite of deductive reasoning. Inductive reasoning makes broad generalizations from specific observations. In inductive inference, we go from the specific to the general. We make many observations, discern a pattern, make a generalization, and infer an explanation or a theory.

Where analytics engines are based on inductive reasoning, the resulting analytics resemble probability theory. Even if all of the premises are true in a statement, inductive reasoning allows for the conclusion to be false. Here's an example: "Harold is a grandfather. Harold is bald. Therefore, all grandfathers are bald." The conclusion does not follow logically from the statements.

This is a better approach than deductive reasoning for projecting the future, but it is obviously imperfect and can produce even more widely varying results.

Advanced IDS/IPS systems use inductive reasoning heuristics to identify malicious behaviors. A heuristic is a rule that provides a shortcut to solving difficult problems and is used when an observer has limited time and/or information to make a decision. Inductive reasoning heuristics lead you to a good decision most of the time, but most of the time is not good enough for advanced threat defense.

Inductive reasoning heuristics are frequently used by contemporary IDS/ IPS systems to generalize the probability of malicious behaviors based on limited input (e.g., known signatures). This also works a high percentage of the time.

Bayesian or Recursive Bayesian Estimation (RBE) Reasoning – This analytic approach is anomaly-oriented and is used in security systems to provide a less tactical view of what's happened over an extended timeframe (e.g. 30 days). Bayesian reasoning is a branch of logic applied to decision making and inferential statistics that deals with probability inference: using the knowledge of prior events to predict future events.

In statistics, "standard deviation" is a measure that is used to quantify the amount of variation or dispersion of a set of data values. A standard deviation close to 0 indicates that the data points tend to be very close to the mean value of the set, while a high standard deviation indicates that the data points are spread out over a wider range of values.

In most Bayesian based security analytics, when a result is 3 standard deviations from normal, the system declares it an "anomaly." The goal of Bayesian Reasoning is to be able to identify a "normal" pattern of behavior by observing subtle fluctuations in activity within the enterprise infrastructure over a period of time to establish a corpus of "prior events". The result is a baseline which is used as a subsequent "benchmark" against which all network activity and/or behaviors will be measured in the future.

Unfortunately, this baselining is flawed and can lead to extraordinary outcomes none of which will result in properly identified threats. There are three significant problems with this approach:

1. If the network and/or the systems being baselined are already infected before the baseline is created then the baseline establishes a false premise,
2. If an insider is already active on a network, the that insider's actions will appear as nominal and become part of the "normal" baseline, and
3. Today's network infrastructure and user behavior is increasingly dynamic, variable and diverse involving many different devices and protocols, access methods and entry points essentially making a baseline assessment impossible without a network lockdown.

Analytics engines that use baselining as their premise for Bayesian Reasoning are prone to extreme volumes of false positives, are cumbersome and difficult to tune and administer, require lots of human attention and frequently miss malicious invasions. In short, they don't work very well.

Abductive Reasoning – Abductive reasoning is a form of logical inference that derives from an observation to a hypothesis that accounts for the observation, seeking to find the simplest and most likely explanation. In abductive reasoning, unlike in deductive or inductive reasoning, the premises do not guarantee the conclusion. This approach is much better suited to the real world of malicious network attacks.

Abductive reasoning typically begins with an incomplete set of observations and proceeds to the likeliest possible explanation for the set. Abductive reasoning yields the kind of daily decision-making that does its best with the information at hand, which often is incomplete.

A medical diagnosis is an application of abductive reasoning: given this set of symptoms, what is the diagnosis that would best explain most of them? Likewise in our jurisprudence systems, when jurors hear evidence in a criminal case, they must consider whether the prosecution or the defense has the best explanation to cover all the points of evidence. While there may be no certainty about their verdict, since there may exist additional evidence that was not admitted in the case, they make their best guess based on what they know.

While inductive reasoning requires that the evidence that might shed light on the subject be fairly complete, whether positive or negative, abductive reasoning is characterized by an incomplete set of observations, either in the evidence, or in the explanation, or both, yet leading to the likeliest possible conclusion.

A patient may be unconscious or fail to report every symptom, for example, resulting in incomplete evidence, or a doctor may arrive at a diagnosis that fails to explain several of the symptoms. Still, he must reach the best diagnosis he can. Probabilistic abductive reasoning is a form of abductive validation, and is used extensively and very successfully in areas where conclusions about possible hypotheses need to be derived, such as for making diagnoses from medical tests, working through the judicial process or predicting the presence of malware.

Securli® uses abductive reasoning to detect malicious behaviors without signatures, baselining and anomaly detection and combined with selected pivot points and contextual algorithms is very successful at identifying malware that has penetrated a network.

Pivoting on Systems Changes Everything

In addition to significant differences in analytic reasoning, Securli®'s other fundamental difference is in its use of the underlying system pivot points instead of event pivots for the analytics.

Most security solutions today focus on events, processing data from billions of events in an attempt to detect malicious behaviors. This approach is extremely limited in its effectiveness; it fails to scale, it generates tons of false positives and noise, and does virtually nothing to parse and reduce the volume of data that depends on automated heuristic analytics in order to produce any actionable information. Events are interesting but taken alone and out of context, yield very little in the way of useful information while creating a large corpus of data points.

The Securli® engine is radically different from other event-based network and/or IDS behavioral analytics approaches because it pivots on systems rather than on events. Pivoting on systems makes activities of the system the focal point instead of events. Our engine seeks out and correlates specific behavioral patterns as part of examining network-based "dialogs" with each system. Each network dialog incident is weighted and once these dialog incidents reach a sufficient evidence threshold within a predetermined diagnosis time window, an infection profile is generated.

Figure 1 - Abductive Reasoning: Infection Profile

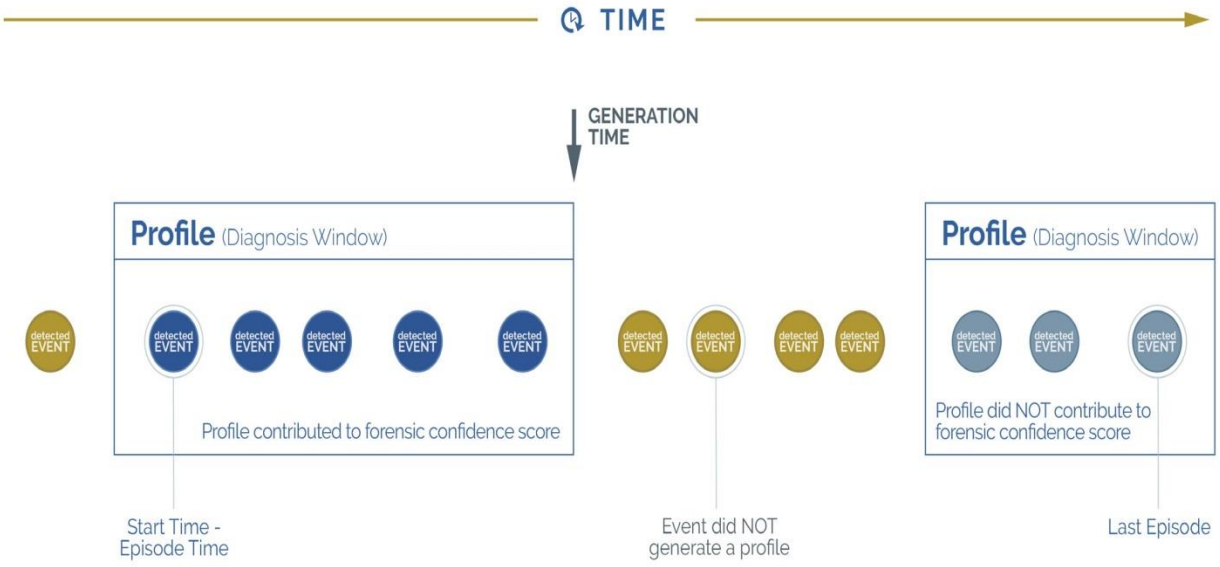

Our engine maps these discovered behaviors to one of 12 stages in the malware lifecycle, each representing different infection profiles. (Figure 1).

As evidence (infection profiles) builds, the likelihood of infection grows and conversely as evidence shrinks the likelihood is reduced accordingly.

Infection profiles are only created against specific network entities (servers and endpoints). This is a crucial distinction since a network may contain 1000's of systems, each generating 1000's of profiles. In comparison, event-based systems like SIEM must process 100's of millions of events on the same network infrastructure. This presents security analysts with three – often overwhelming - challenges: event overload, false positives and lack of context.

Many systems use log data in an attempt to discover and detect malware, but logs by their nature are event driven and today's well-written malware often does not leave a trail in the logs. Traditional log-based data analytics approaches are required to sort through millions of log events in order to correlate those events in any meaningful way. The objective is to somehow convert tens of millions of discrete elements into behavioral patterns.

By pivoting on the events rather than the systems, the data analytics engine must join completely independent variables in an attempt to construct meaningful behavior relationships and therefore must treat every event as potentially significant. This requires tremendous processing power (expensive) and by necessity will generate a very low signal to noise ratio triggering as we have said, many false positives and worse than that, false negatives. In fact, a growing concern with these Bayesian, deductive and inductive reasoning engines is a false negative indicated as an infected system that is determined to be uninfected. To compensate for this tendency, the engines' sensitivity is tuned to err on the side of caution – thus creating even more false positives.

Our engine executes joins on systems as the dependent variables which automatically casts all events in context to the targeted system. This dramatically increases the signal to noise ratio, significantly reducing the chances of presenting false negatives or false positives.

Context is Key

But as we've described, finding the needle in the haystack is only half the challenge. An effective behavioral analytics malware solution also requires context. Network analytics without context is like an orchestra without a conductor - a lot of noise, but nothing that sounds like music.

What we need to do is frame the malicious activities in the context of risk. For example, knowing an exploit is targeting a given system has only limited value. Knowing an exploit is targeting a system that is missing a patch which is making it vulnerable to the exploit has extremely high value. The first indicator is noise, but the indicator in the right context is music to a security analyst's ears.

In our case, context refers both to system context and to the real time nature of the contextual evidence. Effective analytics must track real time activities in order to successfully identify and determine true malicious behaviors. Security solutions that rely on sandboxes to detonate potentially malicious payloads on a state and binary configuration alone are by definition out of date by the time the last system is scanned. Malicious attack vectors change in real-time; they don't remain in state while other vectors are being detonated for analysis.

The Securli® engine correlates real time vulnerability assessments with ongoing threat activity. This involves pulling in Integrity Measurement and Verification (IMV) scans and LDAP (e.g. Active Directory) attributes in real time to deliver the necessary context for security analysts to make decisions.

Our engine provides full interoperability with other vendor products through connectors, inbound REST APIs, and industry standard notation (grammar expressions) for attribution based threat information

exchange within the security community. For additional context, Securli® integrates daily threat intelligence harvested through honeypots, technology partners, and security advisories published by standards-based organizations (e.g. NIST, MITRE, US-CERT).

Unlike traditional IDS approaches which rely on binary (presumed infected or not infected) determinations, Securli® produces evidence-based probabilities of malicious behaviors indicative of malware infection.

Monitored System	Last Episode	Risk Index	Forensic Confidence Score	Threat Classification	External Risk Indicators	Forensic Evidence	Integrated Report (Summary)	Integrated Report (Detailed)
162.229.57.13	04/07/2015 01:52:06 AM		40	Field Alert Malware Scanner Trojan				
162.229.57.11	04/07/2015 01:39:59 AM		38	Field Alert Malware Scanner				
162.229.57.6	04/07/2015 01:48:50 AM		38	Field Alert Malware Scanner Riskware				
162.229.57.3	03/23/2015 12:00:16 AM		38	Worm Field Alert Malware Scanner Riskware				
162.229.57.5	04/07/2015 01:46:08 AM		35	Field Alert				
10.23.212.73	03/16/2015 09:25:17 AM		30	Worm Sasser				
10.23.163.157	03/16/2015 09:25:22 AM		28	External Attacker				
10.23.202.41	04/03/2015 04:07:53 AM		28	External Attacker Botnet Infection				
10.23.179.149	11/05/2014 02:00:35 AM		28	Worm				
192.168.71.136	04/03/2015 04:07:55 AM		28	Spambot				

From a risk management perspective, this is a much more effective and manageable approach because Securli® indicates the systems with the highest probability of malicious infection, rather than presenting binary (infected/not infected) assessments.

As you can see in this example, our Active Risk Dashboard creates a risk index for each monitored system which is delivered continually in real time. This indicates the probability of a system infection and applies a forensic confidence score that coincides with the risk index. It also provides a view into the actual forensic evidence that has been compiled, again in real-time that shows exactly what, who, where and how the risk developed.

Our presentation of the probability that a system is infected (based on the level of malicious activities detected) does require a human analyst's intervention and interpretation, but our brains are much better geared to understand and interpret probabilities than possible certainties, especially when certainty is so rare in cybersecurity. And, the elimination of huge volumes of false positives and negatives reduces the analysts' chore to a manageable activity.

Hiding in Plain Sight

A key differentiation between Securli® and other analytic systems is discerning legitimate from illegitimate callbacks. Callback functions are routines that are passed to Windows API functions as a parameter, and are called later by the API to perform some type of functionality, such as processing messages or handling events. Malware is designed to set up its own malicious callback functions and then pass them to APIs. When the API function is called, the callback function is executed and the malicious code is run.

Today's malware (modern APTs) will obfuscate code, leave no log entries and attempt callback via "trusted" connections to bypass detection by IDS/IPS, SIEM, AMG, sandbox and conventional big data tools. Addressing this new attack vector requires more than just detecting behaviors and generating contextual evidence. It requires monitoring and catching callbacks to a command and control (C&C) server via non-standard channels.

Other solutions map DNS lookups and attempted callbacks to known malicious sites or based on well-known signatures. If there is no knowledge about the site then the callback is considered benign. If the site is a known bad actor, then the callback is considered malicious. It's a binary process.

By comparison, Securli® uses a proprietary Callback Obfuscation and Data Exchange (CODE) protocol in combination with network dialog correlation to detect callbacks based on multiple entropy metrics to introspect behaviors - behaviors that are typically obfuscated. Entropy metrics reflect the lack of predictability of network behaviors. This provides a smoking gun, eliminating the arbitrary inference of other systems.

Even more insidious than obfuscation is a malicious insider. Addressing malicious insiders requires shifting focus from inbound/outbound activities to internal activities.

Another way to look at this is shifting from North-South to East-West (Lateral) movement within the network. Specifically, we need detection mechanisms in place to detect malicious insider activities and the analytics to discern the actual behaviors, their maliciousness and the context in which they operate.

The way most tools are attempting to identify malicious insiders is by using NetFlow data. NetFlow was originally developed as a network protocol by Cisco to help network engineers plan network infrastructure, optimize performance, and better manage traffic and routing. But there are limitations to NetFlow's effectiveness as a data source for analytics.

When a TCP session is established, an audit record (5-tuple, packet volume, duration, AS routing information) is created. When the session is terminated another audit record is created. NetFlow was never meant to be a security tool, partly because of the limited information, but also because everything contained therein is historical in nature. There is no real- time aspect to NetFlow. Malicious insiders operate in and must be apprehended in real time!

To address these NetFlow limitations, Securli® uses a data exchange protocol called SIDE (Signaling Integrity and Data Exchange). SIDE operates in real time and its flow grammar is not static, but stateful. SIDE establishes an initial record and then generates real time update records during the flow. At the end of the flow, SIDE generates a final record.

To establish context, these records contain over 90 attributes beyond basic flow metrics including network addresses, service ports, geolocation indicators, data counters, connection states, threat tags, DNS transactions, connection signals (timeouts, resets), etc.

SIDE provides the basis for visualizing live network flows in real time and in context. This is a paradigm shift, pivoting from network access management to flow entropy management, where the operational integrity of internal and cross-realm network activities and data transfers are now able to be examined. With this approach, Securli® detects lateral data movement, signaling (callbacks, beacons, dial homes), and data exfiltration using flow logic based event correlation in real time.

This opens the door for tactical engagement to intervene in active flows that may be malicious. The threat protocol is extensible and customizable based on policies that are defined using a set of variables, criteria and actions. The policy variables provide attribution about the monitored systems (pivot points), the criteria identify a qualifying episode based on flow variables, and actions prescribe the workflow for automated remediation and incident response.

A set of default global policies are provided out-of- the-box to detect risky (suspect) behaviors and a web portal is provided for security analysts to define local policies based on detailed information about the internal network topology and silos. Polices may also be imported or exported (with the privacy of internal network topology) for sharing the threat definition grammar within the security community.

You are now able to identify a malicious insider behavior (e.g. data transfer from an internal department to a system on the manufacturing floor) in real time to alert dynamic security controls which can terminate the activity before data is stolen, lost or damaged.

Better Data is Better Security

To successfully counter and defeat malware and malicious insiders requires better data analytics, not big data analytics. At a deeper level, we need the right analytics methods using the right detection engines and delivering evidence in real time and within context about the systems we are protecting. This is exactly what Securli® does – it sharply focuses attention on the runtime integrity of assets so it can apprehend malicious behaviors before they transform into a breach.

About Netswitch

Netswitch is one of the world's leading Managed Security Service Providers (MSSP) and the 4th fastest growing MSSP in the world; ranked 5th in California and 61st globally from MSPmentor's 2015 annual top global 501 MSP rankings.

We developed Securli® based on our success in Managed Services as the foundation for changing the way that businesses achieve their IT security goals by providing the most advanced cloud-based solutions to monitor and protect critical information assets without adding headcount or expensive hardware and software licenses.

In business since 2000, with offices in San Francisco, Chicago, Thailand, Beijing, Hong Kong and Shanghai, we provide our customers with experience and expertise in managing their IT infrastructure and defending their networks and applications from cyber-attacks and data breaches.

Small, medium and large companies have all partnered with Netswitch including global clients such as Verizon Wireless, Wells Fargo Bank, Charles Schwab, eBay, Vodafone Americas, Inc., and the Hong Kong & Shanghai Hotels Limited. They work with us because they are confident we have the International reach as well as the local presence to provide both the technology and expertise to complete their mission critical IT projects on time and within an affordable cost structure with integrity and a passion for perfection.

At the end of the day, our customers enjoy the peace of mind they get through knowing we are looking out for them 24x 7×365 days a year.

Netswitch Technology Management

Headquarters:
400 Oyster Point Blvd., Suite 228
South San Francisco, CA 94080

Main: 415-566-6228
Fax: 415-566-4226

Web: www.netswitch.net
Email: info@netswitch.net

San Francisco
Chicago
Hong Kong

"Ranked the 4th Fastest Growing
Managed Security Services Company in the World"

Netswitch™ Technology Management

securli
next generation managed security

For more information, visit www.netswitch.net

Netswitch Technology Management

Headquarters:
400 Oyster Point Blvd., Suite 228
South San Francisco, CA 94080

Main: 415-566-6228
Fax: 415-566-4226

Web: www.netswitch.net
Email: info@netswitch.net

San Francisco
Chicago
Hong Kong