REAL-TIME SECURITY ANALYTICS

A Technical White Paper

NEAL HARTSELL

VICE PRESIDENT, MARKETING

BILLY STANLEY

SENIOR SALES ENGINEER

VICKI IRWIN

DIRECTOR OF RESEARCH, CLICK LABS

JULY 2012

Introduction

In our companion paper, **REAL-TIME SECURITY ANALYTICS, An Overview White Paper**, we discussed the modern security threat, shortcomings of traditional defense-in-depth products to effectively combat these types of threats, an overview of Click's new solution and some example use cases.

This paper briefly revisits the modern security threat, existing product class shortcomings and then focuses on giving the reader strong insight into:

- The Click engine, Click modules and Click Labs
- A set of security analyst use cases
- Real-world case study
- Terminology explanation that sharpens understanding of our solution differentiation

Modern Security Threat Defined

Before we get into security analytics, and more specifically Real-time Security Analytics (RtSA), let's first discuss the modern security threat (MST). We specifically avoid using the term Advanced Persistent Threat (APT), as that term tends to imply that modern threats, breaches, and exfiltration techniques are somehow exotic in and of themselves. In reality, attack techniques have not fundamentally changed in the past few years. MSTs are, in most cases, just automated compilations (but often derivatives that are sufficiently altered to bypass existing products) of previously existing hacker techniques. These techniques tend to fall into a few common 'kill chain' categories: account compromise, beachhead establishment, lateral movement, privilege escalation, staging, and exfiltration. But three things have changed:

First, the attack space is significantly broader these days. IT has been steadily losing its grip on the network environment due to business-empowerment movements like bring your own device (BYOD), virtualization, cloud-based solutions, mobility, the consumerization of IT, and widespread use of social media.

While these trends are driving greater productivity, they expand the attack surface by opening up many new avenues and entry points for hackers to get into networks and perform increasingly nefarious activity. Examples include:

- Employee-owned hardware with weaker security controls, making it more susceptible to infection
- Continued dispersal of corporate data across the cloud and distributed global data centers
- Evolution of application protocols intending to evade traditional security measures, such as onion routing and protocol tunneling
- Growing number of services available inside the firewall
- Increasing number of web applications fronting mission critical apps – while attackers are getting better at web application exploitation
- Employees with access to sensitive information falling prey to hackers' social engineering techniques

Second, while defense-in-depth and the ensuing sprinkling of point products to address specific security needs continue to be de rigueur, IT security staff are forced to manage truckloads of IT telemetry data from products that "show their value" by generating as many events as possible, making it difficult to sift through the noise and find the true source or nature of any attack.

Third, hackers have tooled up with an advanced ecosystem of easily acquired arms and armies of bots to relentlessly stalk every employee, jiggle every door and window latch – and do this non-stop – to your specific network.

Current Solutions are Ineffective against Modern Security Threats

As noted in our first white paper, but which bears repeating here, the typical defense-in-depth armament of today is comprised of four classes of security products – each necessary to broad security protection, but weak at detecting and remediating MSTs:

- In-band signature-based products, like anti-virus and intrusion prevention systems, react to the known bad by pattern matching only on traffic traversing the direct network connection in which they are inserted. They are surgical and fast – but have no peripheral vision and no long-term, stateful memory of contextual attack activity. With tunnel vision and limited contextual knowledge, they stop a declining percentage of the threat types networks confront today.
- Security information and event managers (SIEMs) – despite broader visibility due to their consumption of event and log data from multiple network sources – are 1) constrained by the amount of data that can be ingested and processed in real time (or over any lengthy window of time), 2) only able to act upon what is written to a database post mortem, and only if that data fits a rigid data structure model, 3) constrained by the inherent performance limitations of relational databases, and 4) painfully complex and slow – making it difficult to to build and execute anything beyond the simplest of investigative analytics.
- Policy-based devices like firewalls and identity management products suffer from bit rot and configuration errors – ranging from fat-fingered rule insertion to misclassifications that create new vulnerabilities
- Existing Behavioral and/or anomaly detection products continue to focus on a single angle of attack, such as DoS, DDoS and/or botnet command and control.

Protection against MSTs simply requires much greater contextual and behavioral acuity to distinguish the bad from the good in a timely fashion, and with the accuracy required to allay false positive concerns. MST Attackers have realized that existing security products provide protection against specific attack types and that the key to a successful attack is to continue to vary the type of attack and the specific surface under attack until a weakness is found.

A Fresh Approach to the Detection and Remediation of MSTs

Click Security introduces a solution with a fundamentally different approach to the problem set. Our engine, modules, and intelligence collaboration have fused to deliver the industry's broadest actor-based attribute capture (truly big data), deepest set of analytics (truly big analytics) and fastest ability to crunch these two dimensions against one another – leading

to the breakthrough detection, visibility, accuracy and remediation that enables organizations to reclaim control of their networks.

The Three Solution Pillars

The Click Security solution is built upon three key components:

1. High-performance, memory-based Real-time Stateful Data Flow Engine including the Click O/S and its rich set of interactive visualization and data manipulation capabilities within our Dynamic Workbook
2. Security intelligence encoded into Click Modules
3. World-class security research agency, Click Labs, which drives module development and oversees intelligence collaboration

The Engine

The **Real-time Stateful Data Flow Engine** is the foundation upon which all data collection and security analytic processing is performed. It is instantiated through the Click O/S – which is distributed across Data Mining Units (DMUs) and a Module Processing Unit (MPU). The engine is unique in its ability to waterfall data to a series of interconnected modules – passing only the specific contextual information required to update downstream modules that need it. Through this revolutionary data flow engine, large amounts of telemetry data can be retained in memory for super-fast automated analytics processing.

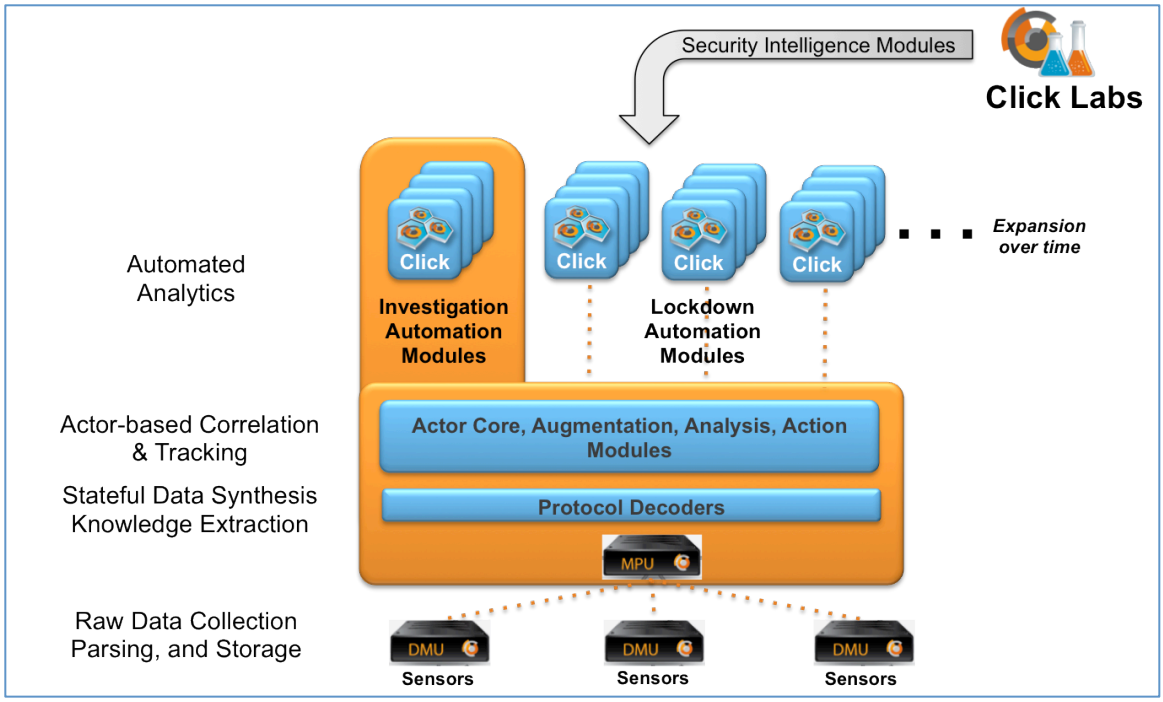

Diagram 1 – Click Deployment

The Modules

Click Modules are programming objects that receive data from predecessor modules, process that data against a security analytic, and perform an output action ranging from 'write results to a downstream click module' to 'invoke a specific human or machine action'. Each module type is described below:

Sensors – The 'shovels' that receive data from telemetry sources, parse and pass the information on to protocol decoders for deeper processing. Sensors are not considered modules, as they largely perform the straightforward work of collecting and forwarding event and log streams to our Module Processing Unit (described in more detail below). Examples of data that is received by sensors, and passed onto Protocol Decoder Modules, include IPS/IDS Events, Netflow Data, Web Server Logs, Firewall Logs, Windows Domain logs, etc.

Protocol Decoder Modules – Essential and complex components of the platform, protocol decoder perform two key tasks:

1. Normalize the various sensor events into flows, authentications, accesses and security events
2. Extract fine contextual details about actors that only decoders have access to, and ensure these fine-grained details get passed to the core, where they will be assigned to appropriate entities. This detailed information about individual actors correlated by entity and presented in a single location is very helpful, and is one of the strengths of Click Security's architecture.

Actor Core Module – Automatically cross-correlate flow, authentication, access and security events – associating each with their respective 'actor'.

Within the core modules, there are four event tables and one entity table. We always start with events, and move, for correlation purposes, to actors. In order to extract actors from events in a general way, we place all varied sensor events on a "common ground", or normalize them. All events collected from sensors collapse to four types of normalized events:

1. ***Flows*** *- these are conversations, simply endpoint-to-endpoint communications. Flows tell us that entity A talked to entity B using such and such protocol.*
2. ***Authentications*** *- these are what you would expect, entity A authenticated to entity B at this time, from that address via this device.*
3. ***Accesses*** *– provides visibility into any type of resource access including web browsing, file-system access, database access, email access, etc.*
4. ***Security Events*** *- these are anomalies and other things normally thought of as security events and are either driven from external sources, e.g., IDS reports, firewall policy violations, click-detected anomalies; or internal sources, e.g., anomalies found through click module intelligence.*

From any of these event types – flows, authentications, accesses and security events – actors are extracted. All actors are collected in the entity table, along with references to each event type in which the actor has been involved. We also keep track of, for each entity, all of the other entities it has talked to. These are called "peers". Fan-out charts are visual

representations of all actor linkages. They are made available through our Dynamic Workbook – the workspace where analysts perform all data analysis and interactive visualization. – Fanout charts are generated by recursively querying the entity table for peers of nodes, and peers of peers of nodes, and peers of peers of peers of nodes, and so on.

Augmentation Modules – Collect and present additional internal and external data, about the actors, that is useful to the investigative analysis. A subset of these modules includes Geo-location, Zone, LDAP, DNS, DHCP, Whois, Web Application Fingerprint, IP Blacklist, and Passive O/S Fingerprint.

Analysis Modules – Analysis modules are the custom analytic that is run against data collected, prepared, and either auto-correlated or made instantly available to analysts for investigative purposes. Here are a few examples of several classes of analytics to give the reader a sense of what we mean:

Rule-based analytics

1. ***Bad IP Range*** – Alerts on successful authentication from within a specific IP range or zone, or from outside a specific IP range or zone.
2. ***Endpoint Restricted Access*** – Determines whether the authentication has granted user access to an incorrect endpoint.
3. ***Resource Restricted Access*** – Determines whether a user or client endpoint has accessed a restricted resource.
4. ***Simultaneous Logins from Distant Locations*** – Alerts on near-simultaneous successful authentications from different/distant geographical locations.
5. ***Brute Force Attempt*** – Detects account logon attempts using different credentials from the same IP within an interval of time.

Statistics-based analytics

1. ***Server Usage*** – *Probabilistic analytic that provides visibility into which departments use a specific server.*
2. ***Primary User*** – *Probabilistic analytic that provides visibility into the primary user of a given server.*
3. ***Server Scope*** – *Probabilistic analytic that provides visibility into the array of servers a specific department uses.*
4. ***Low probability Access*** – *Probabilistic analytic that fires on low probability access from a given department to a given server.*

Action Modules – Action modules are called when an analysis module indicates that an external notification or action is required. A number of actions are possible including:

1. *Email alerting*
2. *SMS messaging*
3. *Automated helpdesk ticket creation and dispatch*
4. *Forensics/evidence gathering (examples: capture full packet data, enable video camera, portscan target to find rogue servers or backdoors)*
5. *Account or badge disablement*
6. *Firewall/IPS reconfiguration (example: add a rule to block particular source or application)*

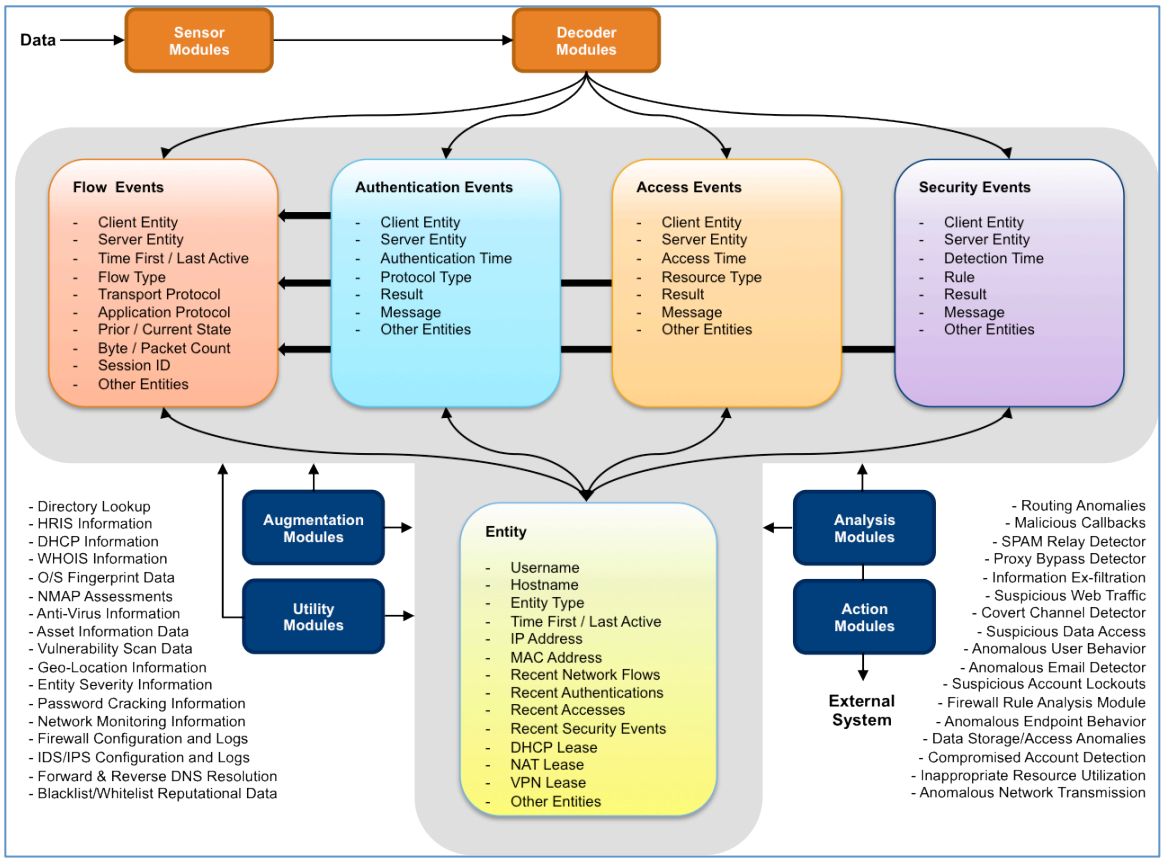

Diagram 2 – Click Module Architecture

The Intelligence Agency

The third component of our solution is **Click Labs**. Click Labs performs three essential functions:

1. Monitors the evolving threat landscape for its own module development effort
2. Assists customers in becoming self-sufficient at writing their own modules as desired
3. Drives security intelligence crowd sourcing through collaboration technology

A growing library of Protocol Decoder, Actor Core, Augmentation, Analysis and Action modules are being researched and developed by Click Labs on a regular basis. It is a vital part of our strategy to keep the solution evergreen with respect to the continuously evolving threat landscape. Our evergreen strategy also includes staying current with respect to external augmentation sources, e.g., IP Blacklists.

However, the system is designed with long-term intelligence leverage and crowd-sourcing in mind as well. Partners and customers can also envision, design and develop modules – either for exclusive or shared use. Modules can be written by tapping raw sensor data, or by leveraging existing Click modules as 'sub-function' building blocks. This flexibility empowers security analysts to capture their hard-earned experience and knowledge

programmatically in a Click Module that can perform 24x7 real-time detection and remediation. The crowd-sourced architecture further enables security professionals to leverage the global community's even broader experience for potentially useful security analytics.

Security Analyst Use Cases

There is much more that could be shared about the power of the engine, the flexibility of the module architecture on top, and the prowess of Click Labs. But to really appreciate what we have built, let's consider a set of occurrences that confront analysts daily – and how Click empowers security analysts in a new way to run these to ground in fast, easy, and extensible manner.

1. Unknown Threat Detection
Zero-day attacks, polymorphic code, and signature-based evasion techniques prevent us from proactively protecting enterprise environments from modern-day malware. The generally accepted norm is that signature-based detections are, at best, 80% effective. But that percentage is deceptive, in that the remaining 20% typically represents a more advanced blend of threats, and more insidious objectives.

Click Security RtSA Response
Our broad perspective of the environment is obtained through a diverse set of telemetry. The core of our product is made up of normalized, elemental data in the form of flow, authentication, access and security information. Each of these events is associated with an entity or actor, which can be a device, user or process capable of performing an action. Identifying and tracking authentication, access, and network flow events at this fundamental level allows us to start building a digital fingerprint of 'normal' on a per actor or resource basis. In doing this, it makes it easy to identify abnormal activity, even though it may not have been fingerprinted as 'malicious' by existing signature-based detections.

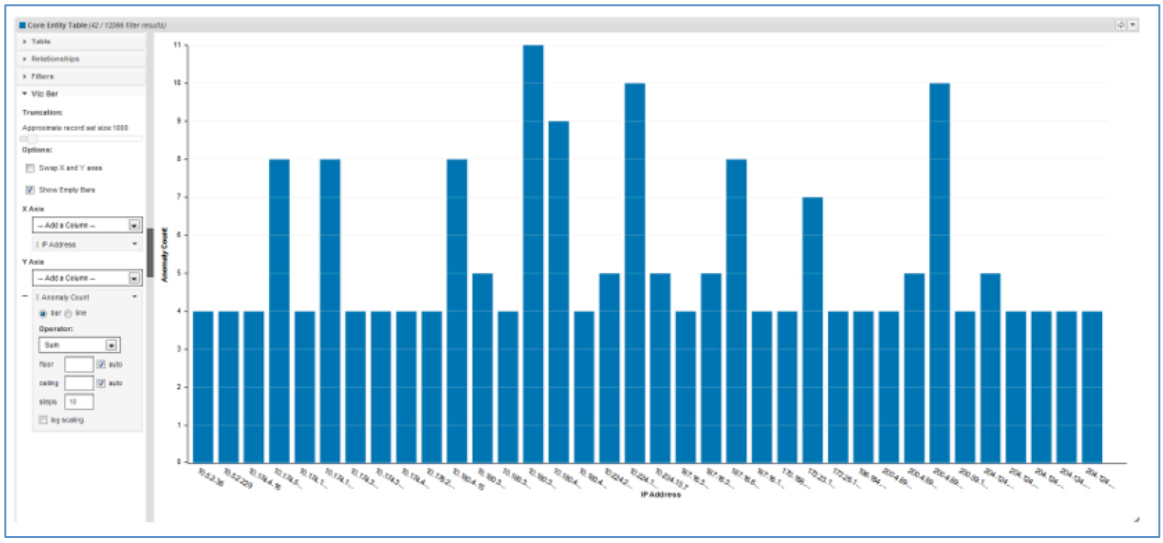

Diagram 3 – Anomaly by Actor Screenshot

As an example, let's consider "George". The fact that George started accessing a private area

on a human resources (HR) file server may not seem malicious. If however, we had enough perspective to know that George worked as an intern in manufacturing, it might seem a little more suspicious. If we knew more about George's 'behavioral fingerprint', we would realize that he has never accessed this HR asset before, and that other characteristics of his authentication, resource accesses and network flows were also behaviorally anomalous. Compounding these insights, there's no doubt that something suspicious is under way. We haven't matched a malicious signature, nor have we observed what we know to be an actual attack. It could be that a compromise has occurred and we are observing reconnaissance. It could also be that an attack is occurring, based on a zero-day threat that is undetectable by signature-based systems. Or, it could be that George has gone rogue and is using the trust and access he's built within the firm to steal intellectual property.

Further, with Click, not only can we profile George; we can also profile the HR asset itself. With an administrator's help, the Click system can identify certain resources as higher value than others. As an example, from the above we can surmise that George is doing something unusual. But George might be new to the company. Or if he works in manufacturing, perhaps he just doesn't use the computer that much, so our "George profile" is still insufficient in and of itself. However, we have developed a strong profile of the HR asset over time. As a result, we know that only Betty and Thelma ever access this particular asset. Now, when we see George poking around in this asset, significant alerts will fire.

2. Embedded Malware Identification
Whether known or unknown, this code has a specific purpose and is designed to be persistent and elusive. Persistence is gained through compromise at various levels within a target system or within various target systems across an organization.

Click Security RtSA Response
Tracking embedded malware typically involves time-consuming, advanced forensics to uncover the actual depth of the compromise. This all assumes, of course, that the compromise could even be detected at some link in the chain in order to spawn further analysis. Deep forensic studies could yield how the malware arrived, when it was executed, what was accomplished, and where it went next. The process of obtaining this data, while time consuming, is also difficult and costly in terms of hardware, software and resourcing. Identifying the depth of compromise in large enterprise environments can take weeks, and that is if everything goes well. By looking at the environment behaviorally, Click makes it easy to understand the path and depth of a complex compromise, conserving valuable time and resources throughout the investigative effort.

As an example, suppose you had a network infected with a variety of malware and the infected systems were communicating over covert command/control channels – either with the mothership, or with each other. If an analyst can simply find *one* infected system, fanout charts should quickly spotlight the remaining machines involved. Here is how. Find one infected system, let's call it A. A phones home to the attacker's machine: Z. One layer of fanout shows A -> Z. But, perhaps there are many infected systems on the protected network all communicating with Z; let's say that B, C, and D are all communicating with Z. Two layers of fanout now give you A -> Z -> [B, C, D]. Next, you just discovered that B, C, and D are also infected. But, perhaps the attacker is also running a server at Y, and C communicated with Y at one time. Increasing the fanout to a third level shows the linkage between C and Y. Increasing the fanout a fourth level shows all the internal systems that

have *ever* communicated with the malicious server Y. Fanouts allow you to keep drilling and drilling until no new linkages surface – resulting in identification of a complete collection of infected systems.

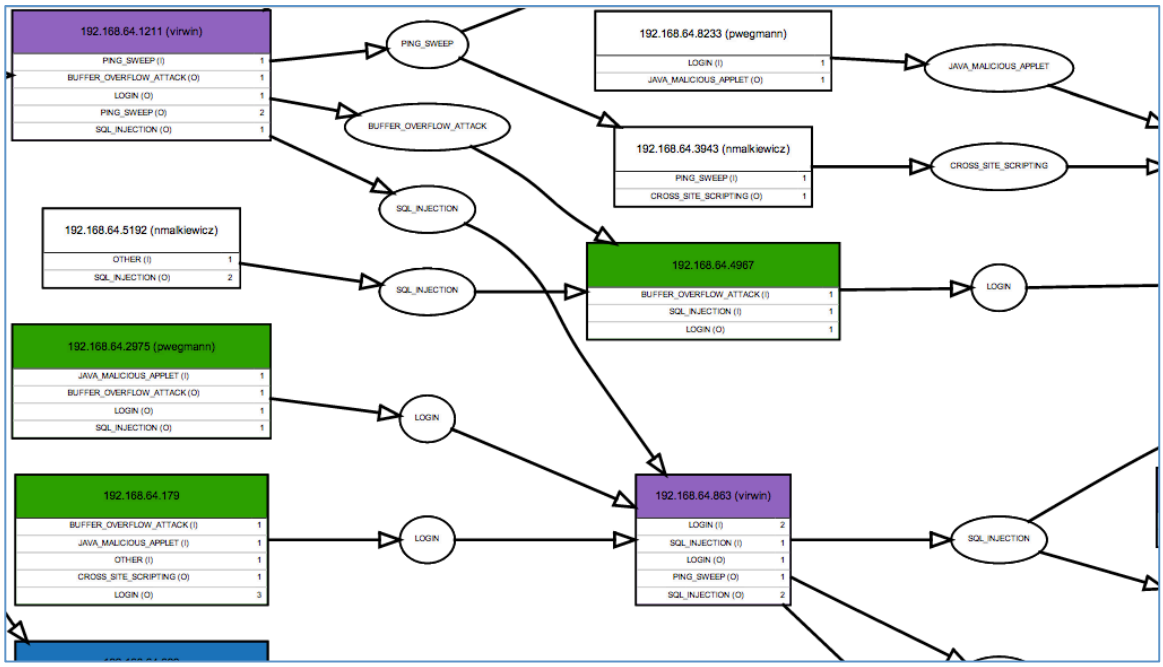

Diagram 4 – Example Portion of a Fanout Diagram

3. Identifying and Prioritizing the Threat

Identifying malware is valuable, but understanding enough about the attack to be able to articulate – or at least hypothesize – the attacker's intent and target is even more valuable. From this intelligence, resource prioritization can be established, in addition to defensive and offensive security postures.

Click Security RtSA Response

Click's solution associates and tracks security events and anomalies as they occur. Both user-defined and anomaly-driven elements are used to identify and prioritize threats, identified from a wide variety of telemetry data. Click's solution highlights analytical statistics that can be used to ascertain and prioritize the threat at the most basic level. Through functional aggregation and correlation, threats can be prioritized with much sharper accuracy and impact assessment.

4. Proactive Security

Becoming proactive begins with detecting what is virtually undetectable with conventional tools. The next step would be to establish an early warning facility that could identify compromise conditions prior to an actual attack being launched. The final step would be to actively engage and neutralize threats proactively, based upon the unknown threat and early warning system.

Click Security RtSA Response

The detection of unknown threats has already been discussed. When looking at the

anatomy of the common attack, we realize there are various early-warning opportunities to engage a threat before an attack is launched. As an example, reconnaissance occurs prior to an attack. This may involve semi-suspicious activities such as network scanning, or it could appear much more innocent, such as network browsing and public resource access. The beauty of analyzing data behaviorally is that the measurement is a deviation from norm rather than a positive fingerprint of evil. For example, profiling allows us to make an educated guess at an entity's "type", based on the kind of events it most often generates. As a case in point, by simply watching kerberos authentications come and go, we can pretty much determine that a particular machine is a domain controller. Similarly, by watching web-browsing activity, we can determine when we have a personal desktop computer at a particular IP address. Then, if we see, for example, outbound IRC traffic from a particular node, even if we don't have a good profile for "normal" for that particular node, the clustering/grouping analysis described above highlights that IRC traffic is a lot more suspicious for a domain controller than it is for a personal desktop. This allows us to recognize early-warning indicators, and neutralize threats before they become viable.

Once the threat has been identified, neutralizing it can be as lenient or harsh as desired. The system can be configured to alert, report, directly engage threats, or feedback new security meta-events – providing a means of having an even higher-level correlation analytic operating on the pre-correlated meta-events). Customers have already invoked system actions including:

- Template-style email generation
- SMS alert distribution
- Help-desk case creation and dispatch
- Active Directory account disablement
- Physical access revocation
- Increased logging levels and issuance of custom instructions to 3rd party systems

Virtually any action is possible with the open-platform underpinning of this solution.

5. Breadth & Depth

Tools today are narrowly focused in terms of their visibility from an architectural and data model perspective. Examples include:

- Authentication information from operating system logs is valuable data that is not necessarily available 'on the wire'. Additionally, certain types of fragmented authentication attacks will be available on the wire, but not always in the operating system logs. Having both sources of telemetry data feeding common, normalized data that can be used for further analysis is most complete.
- Similar to the above, taking the attack down to the link layer by poisoning the ARP cache of unsuspecting victims – suggesting that the MAC of the hacker is associated with the legitimate default gateway address – is a classic man-in-the-middle (MITM) attack. Once the hacker is in the middle of the communication, an authenticated session can be hijacked.
- Numerous hypervisor-style attacks allow attackers to jump between guest VM's and their host, all under the context of the account running the hypervisor, a system-level admin.
- Social vectors enable an attacker to convince a victim to reveal their credentials for physical access controls, ATM withdrawals, etc.

Click Security RtSA Response
Unlike point products – often placed strategically throughout the infrastructure to maintain unilateral perspective – Click gains breadth from broader visibility of a wider range of products across the estate. The idea of depth is also important as you consider that different sources of telemetry data operate at a variety of positions on the stack. The example described above discusses similar authentication-style attacks, launched at different layers of the protocol stack. This is important since point products that are network-based, for example, would not have visibility to application-layer attacks. Likewise, application and network-layer attacks are oblivious to those that are happening at the data link-layer. Click brings each of these elements together, normalizes the data and provides enterprise-wide perspective – a feat we believe no other product can achieve.

6. Data Leakage Monitoring
Managing a corporation's public exposure is essential. There may be situations as seemingly benign as publishing semi-sensitive information externally for easy retrieval at another location, or the intent could be malicious – an intentional ex-filtration of sensitive information. In either case, the result can be detrimental once the information becomes public.

Click Security RtSA Response
Monitoring for data leakage or loss through resource accesses, file system logs, network behaviors and public index services such as Google is a smart way to augment intelligence from traditional data loss prevention (DLP) systems. Modern-day products have limited perspective according to their network placement and traffic visibility. As mentioned previously, Click operates with broad telemetry data delivering a diverse perspective, and at all layers of the protocol stack – offering maximum visibility.

A Real World Case Study

The example security analyst use cases above help to demonstrate the power and flexibility of the Click solution. But, perhaps the best way to understand how our system brings visibility into unknown threats is to walk through a specific breach incident.

Recently, a customer discovered an insidious breach during a product evaluation – where the actor core rollup of interesting activity quickly revealed reconnaissance scanning, possible host misconfigurations, covert channel and backdoor activity as well as active compromises showing internal hosts attacking one another, and in some cases the outside world.

We started by investigating a collection of buffer overflows from an IPS. First, Click fan-out charts revealed attacks from internal hosts were attacking other internal hosts within the customer's environment using techniques that leave no possibility for false alarms – suggesting that the attacking hosts, at least, are compromised. Second, we saw multiple different clients attacking multiple different servers nearly simultaneously and in such a way as to suggest that a single attacker may be controlling multiple systems from behind the scenes. Some hosts were easily identified as being involved in more than one event.

The ability to change perspectives quickly, move up and down the scale between a high

level view (fan-out charts) and a low level view (raw event data) of the attacks – while all the time having the analysis driven by an intelligent anomaly detection and aggregation system, makes it possible to quickly extract important but subtle patterns hidden within an enormous haystack of unsorted, nondescript events. Thus, this capability is a big step forward for security professionals ready to take back their networks.

Detection is the first step, and intrusion detection products do an excellent job at that. But analysis is the second step, and to understand we need to relate events to each other, see how they unfold in real-time, and ultimately translate these events into knowledge about the actors involved. Click Security's solution performs these latter tasks beautifully, providing an unprecedented level of visibility into previously opaque internal communications patterns, and thereby placing power and knowledge back into the hands of the security professional.

Solution / Technology Differentiators

To fully appreciate the power and value of Click's solution, you need to know how our solution is fundamentally different from SIEMs and forensic-based products. That is best handled by just taking a set of terms that we all use frequently and making sure we have the same baseline of understanding:

Actors vs. Events

Events – as well as logs and flows – can be cryptic in and of themselves. Further, you have thousands of events coming every second. In that digital mix, you have four types of events from a security perspective:

- Screamingly obvious that a problem needs attention
- Truly noise and ought to be ignored
- Strong clues, but not screaming – and as a result usually missed since nothing portrays them to you in any context
- Events that, by themselves, yield no clue of trouble. It's not that the clue is vague and accidentally overlooked, it is that the event looks completely benign when viewed out of a unified context.

This is why industry analysts keep telling us 86% of all attacks had evidence of the attack in the event and log streams right under our nose. Click takes events and automatically associates them to actors. Then we automatically associate actors with all other actors. This gives you an incredibly rich contextual picture of what is going on in your network. The blinders come off.

Real-time vs. Hadoop

There is a new fascination with data storage structures based on map reduce technology, and the ensuing ability to perform fast search and retrieval of data once it has been written to that structure, a la Hadoop. No one would dispute that history is important to security analytics. But, we are focused on catching activity as soon as it crosses our doorstep, associating it to actors on the fly, giving you prioritized predictive and behavioral anomaly

alerts immediately. Designing a system to hold loads of actors, actor attributes, and a mass of analytics in memory for real-time, interactive analysis is a fundamentally different problem set than trying to solve for forensic retrieval.

On-line vs. Batch

On-line processing means operating on data as it flows through a set of analytics. It means being able to interact with, analyze, augment and visualize that data in milliseconds to seconds of response time, like a good web search/browse experience. This is what Click's solution is designed around.

Batch processing is what SIEM and forensic tools are designed around. They bring in data, write it to a downstream database, wait for you to pre-decide the question you want to ask, then go about the task of gathering up the records relevant to the specific question you asked, and batch loading them back into memory so you can view them – statically.

The two are diametrically opposed. And, to expect the latter to somehow achieve the former as a 'new feature bolt-on' will leave you disappointed.

Real-time Security Analytic vs. Query

A query, in software parlance, assumes you know exactly what question to ask before you crawl the data. That has four possible outcomes:

1. You are exactly correct in how you framed the query and one search and retrieval exercise will reward you
2. You were close, but not exact, so now you need to formulate a second query
3. You were way off and need to rethink your start point altogether
4. You received the correct answer to the wrong question and walk away with a false sense of enterprise exposure or threat.

The odds of achieving outcome #1 are typically low when trying to find the security anomaly 'needle in the haystack'. So queries into batch processing systems are laborious, time-consuming, frustrating, and therefore expensive.

Click's Real-time Security Analytic turns this problem on its head. First, we use pre-cognitive analytics to piece together a picture of interest ***before*** you asked the first question. **This capability alone gives analysts out-of-the-box value with Click**. Second, we enable you to interact with and refine that picture in real time. Sort, filter, augment, visualize, and fan-out the data until you arrive at exactly an outcome that needs resolution – again in real-time.

Behavioral Analytics vs. Signatures

By now, we are all painfully aware that exact pattern matches of malware are necessary, but wholly insufficient relative to modern attack techniques. But there is a 'hangover' from behavioral security technologies of yore that no one wants to repeat due to oceans of false positives. Network-based anomaly detection (NBAD) and other behavioral-based technologies of the past never gained widespread use for two reasons:

1. Decisions were made on a narrow data set, e.g., netflow data only which led to enormously inaccurate results in the aggregate
2. Only a few analytics – each of which was hard-coded, inflexible, and vendor-imposed – could be brought bear

Click Security has broken both of those paradigms. By ingesting any type of network telemetry data, driving it to actor-level context, and enabling analytics to be flexibly built, on the fly – our solution is capable of running hundreds of sophisticated analytics against millions of memory-based records – in real-time. That nets down to a simple and differentiated value proposition relative to behavioral technologies of the past. We enable the broadest unknown threat detection, with the greatest accuracy, in the least amount of time.

Protocol decoding vs. traditional data collection / normalization

All security products take in some form of network telemetry data, run some type of algorithm, and if conditions are true, take some form of action. And, by extension, taking in data is always parsed and normalized to one degree or another in order to present it to the system's algorithm(s) for processing. That feature set is by no means unique to SIEMs.

SIEM vendors have imposed a rigid data parsing and normalization construct on its users. A construct that then must leave behind important meta data that exists in raw event, log, and flow streams from source products – as it is too difficult to write data decoders or design a back-end database structure that has any hope of decent batch performance if all data is parsed, normalized and stored. Further, most SIEM's are agent-based, and they parse inbound event data in accordance with their own vendor-designed data spec. SIEMs simply use parsing / normalization to do basic aggregation (for performance purposes) and to write elemental data into their fixed database construct.

This is not the same as our concept of decoding. Decoding goes much deeper by deciphering what the events really mean and making an intelligent decision on where the elemental event should land – and which elements of the event will be retained. Unless SIEMs fundamentally change their philosophy and foundation to extract and derive normalization in this fashion, they will never be able to drive the deeper intelligence of which our solution is capable. Additionally, augmenting that intelligence further and performing actions in real-time can only be realized with a deeper intelligence framework.

A good example is our ability to decode windows client-server protocols – which are enormously complex. In that complexity, we are able to maintain the connective particles of state that are lost by simple parsing and normalization. These particles become vital to being able to reconstruct data at the actor level, where data becomes contextual information for the analyst. Many security products track state. But, tracking state across the enterprise and throughout the stack is unique to our solution – and lays the foundation for a new era in data modeling and dynamic threat protection.

Structured vs. Unstructured data

The data within events, logs, and flows can be broadly thought of in three classes:

1. Structured and of obvious value – therefore easy to decide and implement a parse/normalize/store function
2. Structured but of less obvious value – so often discarded on ingest to protect downstream system performance and cost
3. Unstructured – this could be of value or not, but given that it is difficult to parse/normalize/store with in any systematic fashion, the normal decision is just to 'punt' and discard.

Most SIEMs operate off of bucket #1 for ease of design, protection or performance, and avoidance of costs purposes. That is exactly why they are limited in terms of security analytic depth, and therefore behavioral decision accuracy.

Click Security's solution is designed to deal with all network telemetry data equally – structured and unstructured. Clearly, we must impose a structure upon it in order to allow our own analytic modules to operate. But, with the ability to write sensors such that all important information can be retained, enabling ALL data to be post-processed from richer and richer analytics development, we land at a very different place on the cost/performance/accuracy tradeoff diagram.

Interactive Visualization vs. Static graphs / tables

Another unfortunate by-product of traditional systems designed with an underpinning of downstream database storage, batch processing, and structured-data-only visibility is the speed and flexibility at which data can be examined and made useful to analysts. As has been discussed above, legacy technologies – SIEMs and forensic tools – force you to guess the right query (low probability on first attempt), write the query (no easy feat), run the query (if it's more than a simple table join model, it's going to be a while, maybe hours), and finally stare at a statically formed graph or table. What happens when you need to add a field of data, or view the data in a different style? Start all over. Lots of time, energy, and cost wasted. In fact, many analysts tell us it is so horrifically laborious and slow to go this route, they simply give up before starting and just move on to a less strenuous task on the list.

We are purpose-built from the ground up to perform interactive visualization. Trying to achieve what we do out-of-the-box would be analogous to trying to replace the insulation in your walls without going into the sheetrock. The foundation of the product has to be built around the fundamental principles we've built into our core. Without it, trying to replicate what we offer will only add an enormous overhead and performance drag to an already-burdened design.

Click's solution, and the Dynamic Workbook feature in particular, completely upend this model.

Our Dynamic Workbook presents data to you straight from memory, enables you to sort, filter, augment, and visualize that data any way you wish – instantly and with a completely interactive feedback loop. For the first time, analysts are free to run investigative analyses flexibly, repeatedly and rapidly to do what they do best – analyze.

Summary

Modern Security Threats (MSTs) are not novel but they are successful. Big data storage and simple post processing is not the answer.

Click Security's Real-time Security Analytics takes security analytics to a new level using behavioral techniques capable of accurately detecting 'needles in the haystack', linking them into a stronger, more visible picture of suspicion, and providing the analyst with a valuable head start against what appears to be an unknown threat materializing. Whether this is an insider threat or a first strike against a new zero-day vulnerability, Click Security gives you unprecedented visibility and contextual automation. This means stopping MSTs before they progress to the point of loss or damage, not studying them after the fact to hopefully prevent the next one.

But, RtSA cannot be achieved by bolting on a few analytics to an existing product. It must be purpose-built from the ground up – tightly coupling an engine designed for real-time stateful processing of large amounts of memory-based telemetry data with flexible; easy to write actor-based analytics; interactive visibility for specific investigation work and broader exploratory work; and an ability to capture and share crowd-sourced security intelligence on the fly.

Click Security's solution and technology affords security analysts a massive head start in seeing and understanding the unknown BEFORE it progresses to the state of loss/damage. The value of prescient visibility cannot be overstated. Indeed, as is well chronicled in one of the industry's leading bellwethers – **Verizon's 2012 DATA BREACH INVESTIGATIONS REPORT** – considering how long it takes organizations today to even realize they've been compromised, the early, visible detection Click Security can provide could dramatically reduce breach incident loss/damage.