Predictable Security

A next generation

PREDICTIVE

solution to cyber security that *every* organisation should review

Presented by **Computed Future**

Document version 0.1

Introduction

Today's organisations are besieged by security threats. Like kids in a candy store, cybercriminals and nation states can't wait to get their hands on confidential information for gain or to wreak havoc. As attacks grow more advanced, it's increasingly important for organisations to have more sophisticated security tools in place to meet changing security requirements.

Predictable Security is a revolutionary solution to the cloud security problem. In this white paper, we introduce the high-level components that form part of cloud computing environments. From there, we show how Predictable Security addresses the challenges of protecting these components. Finally, we will clearly demonstrate Predictable Security's value proposition and how it differs from the current vendors that play in this market.

Predictable Security is the only product on the market today that can predict future attacks - providing a revolutionary solution to the cloud security problem.

The problem

The most widely used definition of the cloud computing model is introduced by NIST as "a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction". The biggest motivator for the adoption of the cloud is its ability to reduce IT costs and increase capabilities and reachability of the delivered services. Despite the benefits, the approach of cloud computing still has a lot of open issues that impacts the model credibility and pervasiveness.

Security is the major concern that hampers the adoption of the cloud computing model because:

- Security management is typically outsourced by enterprises to a third party that hosts their IT assets, thus creating a problem of loss of control.
- Different tenants sharing the same assets in the same location and using the same instance of the service while being unaware of the strength of security controls used.
- The lack of security guarantees in the SLAs between the cloud consumers and cloud providers. Hosting this set of valuable assets on publicly available infrastructure increases the probability of attacks.

The problem can be practically illustrated by a real life example: Imagine you purchase a bar of gold as an investment. You have done a lot of research and the price of gold is rising - providing a very lucrative opportunity. You purchase a state of the art, fireproof safe which you build into your home using steel and concrete. You also go as far as installing a sophisticated alarm system connected to a security firm to provide an alert should the safe be breached. The next day you go to your favourite holiday destination as you are expecting to make a good amount of money with the gold investment.

While you are enjoying some well deserved R&R, you get a call from the security company monitoring your alarm. They inform you that over the past few days, burglars had broken into your house through a window. They are not sure when, but the attackers went to work on cracking the safe. Once they succeeded, the alarm was sounded, but the gold bar was stolen before the response unit could reach your house. You took every precaution to protect your gold bar, but it still wasn't enough.

This example mimics what could happen on your cloud infrastructure. Firewalls etc. prevent direct access to your network but invariably there are opportunities to gain unlawful access. Once inside, malicious operators take their time breaking into your secure entities. IDS systems may pick up that an attack is happening or that something bypassed your firewall or IPS but it may be too late for evasive action. These attacks often go undetected due to the fact that all current approaches and technology available, only prevents you from known attacks or threats or once it has already happened i.e. protecting you from the past.

Now, imagine that you had received an alert as soon as events suggesting burglars opening a window and stepping inside was detected and then a more severe alert once they started tampering with the safe. Both these alerts would have enabled the security company to come out and check on your asset, foil the attack and enable further action to secure. Now go a step further and imagine that the security company could inform you

two days before the breach is due to happen, that an attack is potentially imminent (based on the fact that it has noticed a behaviour at the front of your property, including similar people/faces observing your house for a number of days, plus a vehicle with the same number plate pulling up at the front of your property) and that you should act. This ability to predict an attack and the possible outcome of an attack is the power and unique differentiator of Predictable Security - it detects attacks early and provide ongoing monitoring of severity - enabling fine grained action to diffuse the risk.

Cloud security

We summarise the key security issues/vulnerabilities in each service delivery model. Some of these issues are the responsibility of cloud providers while others are the responsibility of cloud consumers.

VM security

Securing the VM operating systems and workloads from common security threats that affect traditional physical servers, such as malware and viruses, using traditional or cloud-oriented security solutions. This issue falls under the responsibility of the consumer.

Securing Docker image

As part of cloud consumer responsibilities, developers need to make sure they are downloading Docker images from trusted sources that are curated by the Docker community or the vendor and run vulnerability scans against those images before running them in the host environment.

In the breakdown of process or if the process does not exist internal to the consumer, the responsibility should revert to their security solution. Predictable Security is one such solution that will generate an alert if unverified Docker images are downloaded from untrusted sources

Unsecured communication and unrestricted network traffic

In some versions of Docker, all network traffic is allowed between containers on the same host. This increases the risk of unintended and unwanted disclosure of information to other containers.

Predictable Security will generate an alert for every connection or packet exchange that is not linked between specific containers and will generate alerts for communication that does not have TLS enabled when it communicates with docker registries.

An example test case

When a container accesses a database or service, it will likely require a secret, such as an API key or username and password. An attacker who can get access to this secret will also have access to the service. This problem becomes more acute in a microservice architecture in which containers are constantly stopping and starting, as compared to an architecture with small numbers of long-lived VMs.

Securing VM boundaries

VMs have virtual boundaries compared to physical servers. VMs that co-exist on the same physical server share the same CPU, Memory, I/O, NIC, and others (i.e. there is no physical isolation among VM resources). Securing VM boundaries is the responsibility of the cloud provider.

Unrestricted access of processes and files

If the attacker has root access to the container, they may have the ability to gain root access to the host, often through vulnerabilities in the application code. Access control best practice recommendations include the principle of least privilege. The user namespace feature in Linux containers will allow developers to avoid root access by giving isolated containers separate user accounts, and mandate resource constraints, so users from one container do not have the capability to access other containers or exhaust all resources on the host.

Predictable Security will generate an alert if the user has root access to the container and generate an alert if the user with root access to the container tries to read/write access to any directory on the host using any method.

An example test case

An attacker who gains access to a container should not be able to access other containers or the host.Users are namespaces by default, so any process that breaks out of the container will have the same privileges on the host as it did in the container; if you were root in the container, you will be root on the host. This also means that you must be on alert and concerned with potential privilege escalation attacks whereby a user gains elevated privileges such as those of the root user. This is often accomplished through a bug in application code that needs to run with extra privileges. Given that container technology is still in its infancy and although container breakouts are unlikely - just like any other breach - care should be taken to ensure you have adequate protection in the event that such an attack occurs. Predictable Security will offer protection in the event that an attacker does obtain root access due to overlooked vulnerability.

THREAT CATEGORY	VULNERABILITIES	ATTACK EXAMPLES
Vulnerable Systems and APIs	Hypervisor bugs Unpatched software.	CVE-2017-10912-Guest to host Privilege access
Denial of Service Attacks	Flawed network architecture Insecure network protocol.	Memcached attack to create the biggest DDOS attack seen on the world. Github Attack.
Shared Tennant Vulnerabilities	Virtual Machine Vulnerabilities Hypervisor Vulnerabilities Hardware Vulnerabilities	Spectre and Meltdown Attacks

Approaches employed by current vendors

Agent based

Depending upon the type of workload (server, client) and type of platform (Windows, Linux) endpoint security is deployed. The approach is no different than endpoint security agents (typically anti-virus, next generation anti-malware) deployed on a host. This approach relies heavily upon using signatures in order to detect threats that are well suited to windows and some linux malware attacks. The challenge with this approach is that cloud workloads - especially container attacks are novel and it's almost impossible to detect these attacks with traditional signature approaches.

Network based (inline)

The inline network based security approach is mostly used in cloud security as a service. Depending upon the type of deployment to access cloud security as a service, the entire access to the internet is relayed/routed

through the security controls that cloud security as a service is providing. This approach doesn't take the cloud provider into account. The entire approach is independent of the cloud provider. Hence, any attack that targets the cloud infrastructure won't be able to be mitigated by this approach.

Limitations of current approaches

The major drawbacks would be that the above approaches normally send too many alerts to monitor, they are based on signatures that fire on successfully executed attacks and they cannot predict attacks that are underway. Further to that, these types of defenses are also visible to virtual machines or containers so attackers can profile them and subsequently change their behaviour.

How we do it differently

Overview of state of the art in machine learning in security

The security industry has jumped on the Machine Learning bandwagon. This is mostly due to the competitive pressures, forcing them to claim Machine Learning capabilities after a recent surge in hype on this technology. The reason for this surge is the advancements achieved in certain machine learning algorithms, collectively called "deep learning". These algorithms are mostly based on refinements and improvements of neural networks. Neural networks have been researched since the 1970s. During the so-called "AI Winter", research into neural networks and machine learning in general was relegated to a smaller group of researchers that held onto the initial promise of neural networks. The tenacity of this smaller group and the growth in computing capacity led to significant breakthroughs in the last decade.

A hype cycle ensued and every prominent technology company, including security vendors, now nearly claim to do machine learning – mostly based on these new incarnations of neural networks and related techniques.

What is seen in the security market today is the following:

- Supervised machine learning to detect malware, detected spam and phishing content. Examples of vendors that use supervised machine learning include Cylance, Blue Vector.
- Anomaly detection applied to breach detection, fraud detection and impending system failure.
- Unsupervised machine learning to do forensics and replace manual rule-based pattern matching, usually via manually executed data science projects.

Almost all the above is done by the extremely popular deep learning algorithms developed in recent years. These algorithms rely on differentiable inputs and parameters to find the best fitting model that can classify numbers into likely categories or predict an outcome based on some input numbers. Unsupervised algorithms and anomaly detection algorithms also require the application of numeric processing on parameters.

Our unique approach to machine learning

Over the past 10 years we have been working on a revolutionary approach to machine learning that does not suffer from the issues listed above. This technique is based on research that Fritz Venter started in 1993. The research technique has also been the reason for the US granting him a green card within 1 year from the date of application based on a special class of petition that some call the "genius visa". Fritz has been published in various journals and conference proceedings (listed further on in the document).

Fritz is the inventor of various patents that are in progress of being issued and to date, he is the inventor or co-inventor of various issued US patents (listed further on in the document).

During this time we developed the core machine learning training and inference algorithms based on the outcomes of our research.

Another one of our co-founders, Dr Bruce Watson has published numerous papers on this subject. Bruce's publications can be accessed at ResearchGate (link).

We also filed a patent protecting the technology underlying Predictable Security. This patent contains 36 claims, including 7 primary claims. A follow-on filing that is underway will increase the number of claims significantly.

As an example of our unique graph-based machine learning approach, the below 2 figures are screen shots from our user interface:

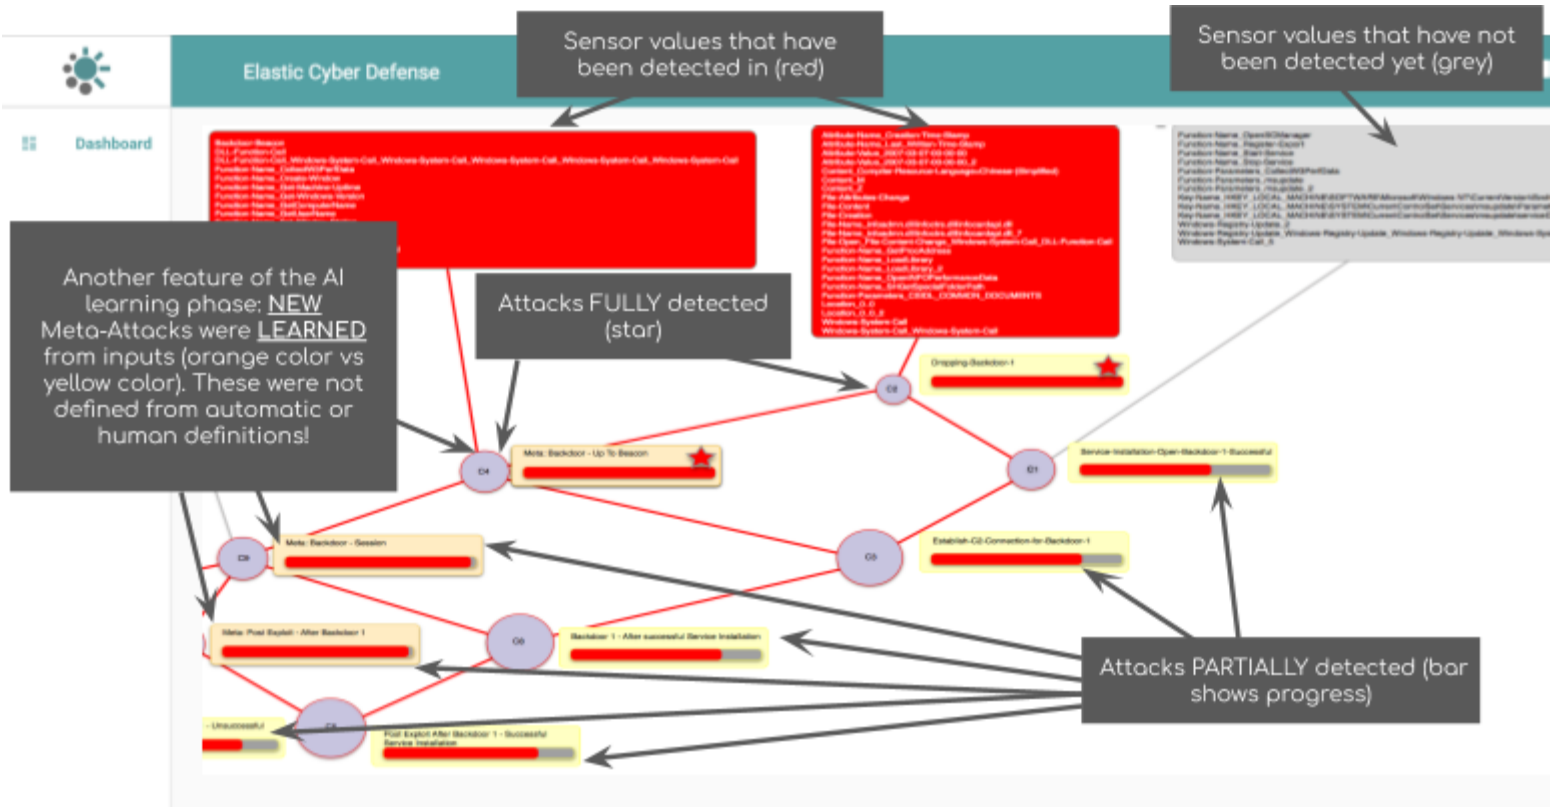

Figure 1: Screenshot of Predictable Security in action

In Figure 1 we show the progress of an attack by the well known Deep Panda APT.

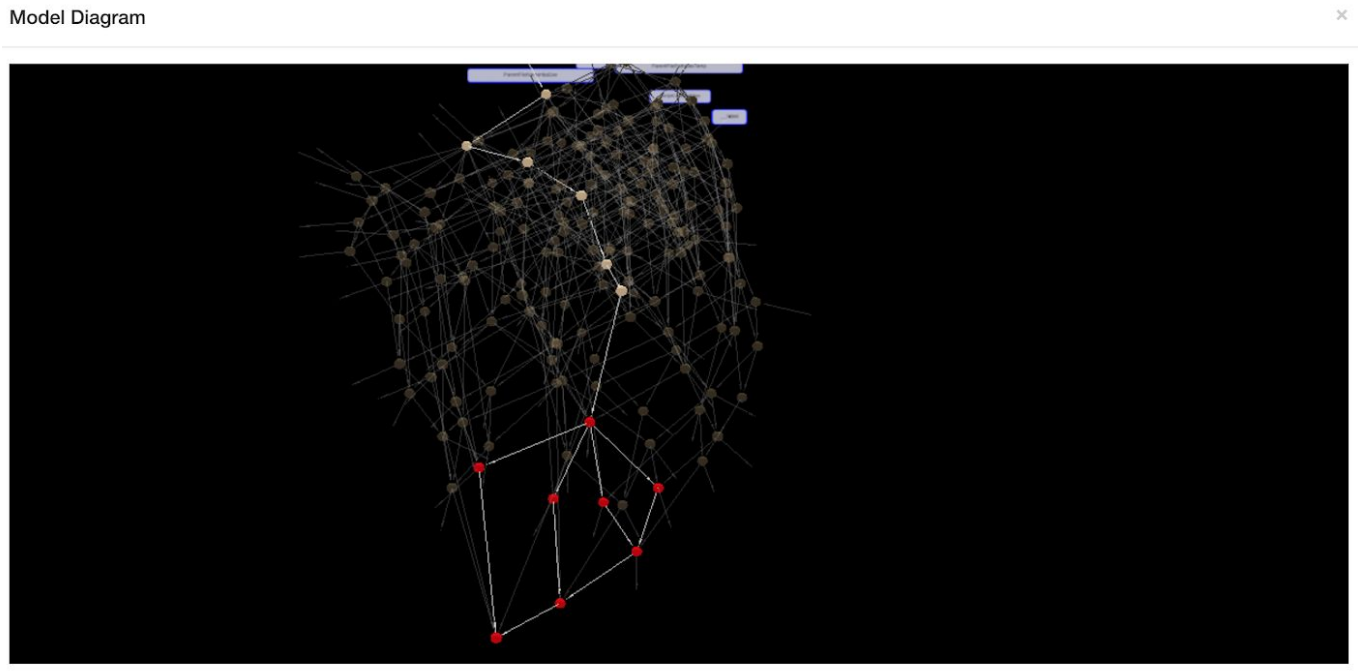

Figure 2: Progress of an attack

Figure 2 depicts the progress of an attack during inference (off-white nodes) and the fact that all possible outcomes from here are malicious (red nodes).

PCS | Events | Models

Attack Progress	Model	Session Originator	Session Id
0	EndpointProcesses	Attribute_Feeder_VM_1	47b46e7d-bcf3-49c6-afba-97a990c4f2b5
0.7142857142857143	EndpointProcesses	Attribute_Feeder_VM_10	30aac1b5-e283-438e-9d32-8a7b1db1b61c
0	EndpointProcesses	Attribute_Feeder_VM_12	18b682af-91a2-4b49-9e27-75a1c58f7a58
1	EndpointProcesses	Attribute_Feeder_VM_23	41343d1d-e6db-4619-87ca-bd12528022bc

Figure 3: Screenshot of inference sessions underway

Figure 3 shows a small snap-shot of inference sessions that are underway. White sessions are still far from an attack starting. Off-white sessions are predicting possible attacks underway. Red sessions are successful attacks. The attack progress of all running sessions form the basis of predictive alerts that the system forwards to the customer's SIEM.

Our approach is completely different from the in-vogue neural network/numeric processing family of machine learning techniques in use by security vendors today for the following reasons:

Discrete/Symbolic Machine Learning

Our approach is in a completely different branch of the general machine learning field that concerns the world of discrete features or events. To handle sets of discrete categories in neural networks, discussed above, categorical inputs need to be encoded into numbers so that neural networks can process it. Our approach handles discrete sets of features directly.

Unified Supervised and Unsupervised Machine Learning

Our approach is differentiated from numeric and other symbolic machine learning techniques in the following way: We can capture all inter-feature-set relationships, not just the relationship between a predefined set of inputs and a target output as is the case in supervised machine learning based on other symbolic machine learning techniques. This also means that a model trained using our approach handles both supervised and unsupervised machine learning. It also means that our approach generalises to multiple use cases much easier than deep learning or other symbolic machine learning approaches. The practical value of this characteristic of our approach in the security field is that our approach can model all relationships prevalent in a multi-dimensional space of system states captured from diverse sources, including hosts, virtual machines, networks, user behavior, etc.

Understand sequences of events from sensors widely distributed in space and time

Another important differentiator for our approach is that we can model sequences of system states. This makes it possible to model the relationship between sequences of system events.

We do not simply classify a snapshot of a system, a binary file, dynamic or static content as is the practice in current applications of machine learning in security. We can create models that predict possible future sequences of events that lead to malicious outcomes before they occur.

Glass Box vs Black Box: Explain attack history and future

One of the well-known drawbacks of neural networks is that they are not able to explain why they predict a specific classification or value in the case of supervised machine learning. Similarly, current numerical un-supervised machine learning techniques cannot explain why certain data points cluster together. This is also why a model trained using numerical machine learning approaches is sometimes called a "black box".

In contrast to the "black box" nature of numerical approaches, our "glass box" model is able to explain its assessment of historical sequences of events and explain its predictions of possible future sequences of events. This transparent view of history and the future of state as it relates to time can be visualised on the system console or communicated to a customer's SIEM using log forwarding.

Fine Grained Recovery: Roll back only malicious sub-steps

A very important additional value of our approach is that we can roll back state of any system by "walking back" the historical event sequences that a model has tracked up to a pre-malicious state and apply the state snapshot of any system associated with such a pre- malicious state.

Out-of-VM/Out-of-container security

The fundamental concept with regards to providing out-of-VM/out-of-container security is to leverage the hypervisor to provide better security. Tapping into the hypervisor provides following advantages:

- Better context – It provides protection from outside the guest OS, and provides the protection capabilities from a trusted context .
- Tapping into the hypervisor provides new capabilities – it helps to view all interactions and contexts of:
    - CPU;
    - Memory;
    - Network; and
    - Storage.
- More cost-effective security - instead of deploying hundreds of agents, across different VMs, deploying security control on hypervisor gives a single pane of glass and helps understand the context and interactions across hundreds of guest operating systems.
- Overcoming limited visibility into the host OS (vs. in-VM approaches) and virtual network to find vulnerabilities and assess correct configuration.
- Alternative tools required to do these are virtually non-existent and are of limited in scope targeting only a few platforms.

How it works

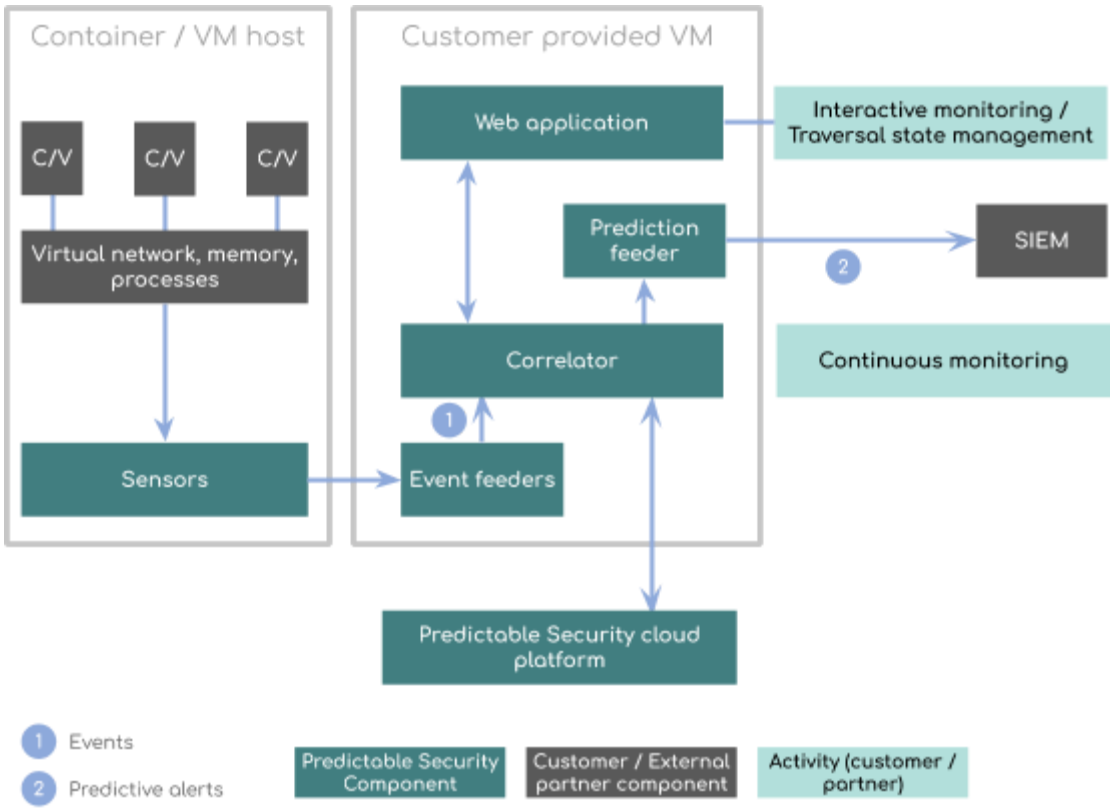

Figure 4: Predictable Security logical architecture

Overview

The below section will provide an overview of the logical architecture for Predictable Security.

Definitions

COMPONENT	DESCRIPTION
Container/VM hosts	Refers to where we install our low level sensing technologies with minimal invasion into customer environment and minimal impact on performance of customer computing capability.
C/V	Refers to containers/virtual machines where we sense low level data from.
Virtual Network	Refers to the virtual network where we sense low level data from the virtual network traffic, memory and processes of containers and virtual machines.
IDS/Sensors	We provide our own sensors to extract low level data from monitored resources such as containers or virtual machines. We can also use outputs from standard signature based IDSs that the customer already uses.
Event feeders	Event feeders are installed on a customer provided virtual machine. They are responsible to extract higher level events from lower level events.
Correlator	The correlator is also installed on a customer provided virtual machine. It consumes events sent by event feeders and produces information about potential attacks in

	progress and predicted further attack steps. It makes use of pre-trained predictive models that are continuously updated by our cloud platform.
Prediction feeder	This prediction feeder is also installed on a customer provided virtual machine. It feeds predictions produced by the correlator to the customer's SIEM. We currently support syslog compliant SIEMs and will continue to build out our list of supported SIEMs available upon request.
Web application	The web application is an optional application also installed on a customer provided virtual machine. It provides a web based monitoring tool of the same predictions generated by the correlator and sent to the customer's SIEM. It helps to also verify the installation of our components on the customer environment.
SIEM	SIEM stands for Security Information and Event Management system. Most customers use a SIEM to monitor their systems from a security standpoint. We currently support syslog compliant SIEMs and will continue to build out our list of supported SIEMs available upon request.
Continuous monitoring	Continuous monitoring of the predictions sorted by attack progress (from furthest progressed to least progressed). Information on why predictions are made based on seen event sequences are provided as part of the continuous monitoring.
Customer provided VM	Two or more virtual machines (depending on the volume of data flowing through the system) where our components are installed and run on.
Interactive monitoring / traversal state management	On the optional web application you can monitor predictions from all event originators and drill down to see why we are making the predictions we are making. We can also clear sessions and perform management of the correlators internal state (e.g. downtime for memory, disk etc. upgrades).
Predictable Security cloud platform	This is the central platform hosted by us that is used to: • Train new models based on new threats or vulnerabilities; • Re-train existing models as new training data is collected; • Continuously update customer models used by the correlator; • Continuously update our software components when new releases become available; • Remote install our product components on a new customer environment or to protect existing and new hypervisors/hosts in the event that a customer adds new ones or more resources; and • Collect usage metrics. Our cloud platform will form the basis for a future hosted user interface, SIEM and monitoring service for those customers that do not want to use their own SIEM.

Step-by-step

Installation of the Predictable Security suite of components is a simple process that can be done remotely after certain prerequisites have been met for the customer's environment.

Our sensors run on the host that they are installed on and thus not another burden, overhead or more load for each individual container (also results in less installations and subsequent upkeep). Regardless of whether the traffic is encrypted or not, we scan for patterns that appear or that is picked up based on known events and more importantly, ***unknown events.***

These events are then sent to our correlators (installed on a virtual machine provided by the customer) where the events are then sent for some magic:

- Based on our secret magic sauce, we run the event or series of events through our patent pending machine learning algorithm and calculate where you are in the threat cycle (in the event of a malicious event/breach/attack) and based on the seriousness, provide an appropriate alert;

- As part of the predictive alerts that we send to your SIEM, we also provide information about the event sequence that we base our predictions on (we support the main stream syslog integration standard SIEMs and will continue to build out our list of supported SIEMs available upon request);
- We keep the correlator installed in your environment up to date by continuously training existing and new threat or vulnerability models and updating the models deployed for your monitored resources.

FAQ

Is the product an IDS?

The product is not an intrusion detection system but a sophisticated intrusion prediction system that is based on a unique patent pending symbolic machine learning approach.

Is the product an IPS?

The product is not an intrusion prevention system but a sophisticated intrusion prediction system that is based on a unique patent pending symbolic machine learning approach.

How does it differ from an IDS?

An IDS is reliant on known events in order to identify a threat or attack. In order to determine an event that is a threat or attack (bad event), an IDS relies on signatures that are based on known threats or attacks (bad event). Predictable Security relies on our patent pending predictive capability to determine if the event that is in progress or being processed will lead to an attack or a bad outcome.

Will it replace my IDS?

Predictable Security can act as a compliment to your existing IDS adding that predictive capability to your environment . As an example, Predictable Security has the ability to consume alerts from your existing IDS and further enhance its capability.

How does it deal with zero day vulnerabilities?

It provides predictive capability to determine if the current event will lead to an attack or bad outcome. If you take into consideration that an attacker can linger around for an average of 180 days, meaning that zero day is actually not day zero but day 180. Predictable Security will monitor all events, essentially determining all activities that could lead to a bad outcome thus possibly detecting 0 day vulnerabilities before they happen.

How will it integrate with my SIEM?

We can feed the predictions to syslog-compliant forwarders that integrate with most major SIEM. We can also integrate into any SIEM that provides an industry standard interface such as REST and will continue to build out our list of supported SIEMs available upon request.

I don't have a SIEM, how will I be alerted?

We have 3 options:

1. We can install a SIEM for you;
2. We can send alerts to an email address; or
3. You can view attacks-in-progress on our user interface.

Is it visible to attackers?

No, we run our sensors in the hypervisor or the container host.

Does it require a lengthy installation process?

No, we can install the system remotely or with minimall on-site effort pending organisation size, complexity, SIEM integration and customisations.

How does it do the prediction?

It runs inference on every event streamed into the system. Inference is essentially a set of patent pending traversal operations on the underlying graph representing the trained model. For example:

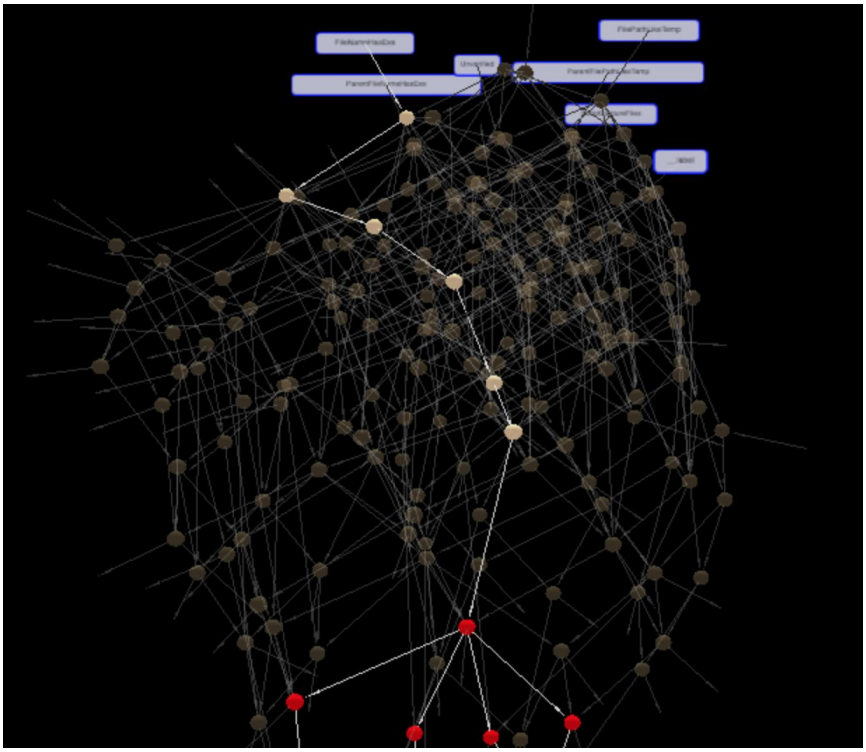

Figure 5: Screenshot showing how inference is tracking

In Figure 5 we show how inference is tracking a specific "path" in the graph and the nodes below the tracked path are all possible paths from specific node. These nodes are used to calculate the prediction of future events. In this case they are all red, meaning malicious nodes. Nodes that are faded have been removed from the multi-dimensional space of possible event sequences.

Why is this a good way to make predictions?

This is a good way to make predictions because the algorithm does this very efficiently in comparison to other machine learning techniques such as neural networks.

How much detail does it provide with regards to the prediction?

We provide progress as a number between 0 and 1, the steps that got us to the current progress and possible future events that may follow as well as how that brings us closer to possible successful attacks. We also provide the Monitored Resource (for example the VM number or container name) that the prediction is for.

How does it protect my cloud workload?

The system protects your cloud workload by consuming low low level data from virtual networks, container hosts and hypervisors that run any containers or VMs that runs the customer's work loads. This stream of low level data is then transformed and run thru our correlators to do the necessary inference and stream out the predictions. We can therefore monitor and protect any part of a customer's workload on the cloud.

How does it work with my existing CI/CD pipeline?

As soon as you run your CI/CD pipeline you may download base images from external repositories that may introduce compromised containers. The system will see any change in behavior of these containers as soon as the external repositories are accessed from your container host or when the updated containers start running.

What happens when we run a pen test?

In the event of a simulated penetration test, Predictable Security will generate alerts as expected. We can work with our customers to limit the alerts that are sent to their SIEM in the event of a pen test.

Final words

We believe that Predictable Security is the only solution capable of moving organisations ahead of attackers. Every enterprise/government entity should at least review the capabilities of Predictable Security and asses its value within their current environment.

When cloud computing started, the idea was to move the workload from an on-premise model where we were responsible for managing the life cycle of our infrastructure to a model where the infrastructure management and operations are run by somebody else at a much greater scale. The general idea was to reduce the computational cost, operational costs and to provide easy scalability for future needs.

This overall premise didn't take into account the security of the infrastructure where cloud workloads were involved. As time passed, different attack variants targeting cloud infrastructure as well as workload emerged that couldn't be solved by traditional security measures. Moreover this infrastructure has unique security issues for technologies such as containers, CI/CD (continuous integration/continuous delivery) pipelines, orchestration solutions e.g. kubernetes that traditional IDS/IPS/HIPS are woefully inadequate to address.

To further complicate the situation, based on our research, attacks that didn't work on physical machines/VM due to security solutions (HIPS/IDS/IPS) have since become invisible to these security solutions if the attacks are performed on hosts that are containerised. An example of this would be CVE-2016-5195 which can be detected by most of IDS/IPS/HIPS when run against a physical linux kernel/VM, these solutions are blind when this vulnerability is leveraged inside the container to do a privilege escalation.

This gives us a unique opportunity to look at the artefacts inside these containers from a network as well as a system perspective. We feed all events that occur to our Predictable Security engine to then perform an in-depth analysis (our secret magic sauce) on these potential attacks and alert you to the outcome. Predictable Security is a groundbreaking product due to its ability to assess the critical path of an event/potential attack and then ***predict*** the possible attack paths and outcomes in order to alert you to previously undetectable events and possible attacks.

About the authors

Fritz Venter

Fritz's career spans 24 years in various technical and leadership roles. He has published numerous papers on Formal Concept Analysis, Machine Learning, Pattern Matching and Knowledge Discovery in Databases. He is an inventor of various USPTO registered patents in advanced analytics. Fritz has also been very active in building startups, managing engineering teams and applied machine learning research. He has a Masters degree in Computer Science.

- Modelling the sensory space of varietal wines: Mining of large, unstructured text data and visualisation of style patterns (link);
- Images & videos: really big data (link);

- Pattern Matching using Position Encoded Pattern Lattices (The 9th International Conference on Concept Lattices and Their Applications, Springer October 1, 2012);
- Failure Deterministic Finite Automata (Prague Stringology Conference 2012, Prague Stringology Club August 1, 2012);
- Pattern Matching in Structured Multi-Sensor/Layered Image Big-Data (Presented at: 33rd Canadian Symposium on Remote Sensing July 1, 2012);
- FCA-Based Two Dimensional Pattern Matching (LNAI Formal Concept Analysis: 7th International Conference, ICFCA 2009 Darmstadt, Springer-Verlag Berlin, Heidelberg May 1, 2009);
- Knowledge discovery in databases using lattices (Expert Systems with Applications, Elsevier November 1, 1997);
- Lattice-based Knowledge Discovery in Network Management Data (Informatica an international Journal of Computing and Informatics, Slovenian Society Informatika January 1, 1997); and
- Using a lattice for visual analysis of categorical data (IFIP series on computer graphics : Perceptual issues in visualization, Springer Berlin ; New York January 1, 1995).

Fritz is the inventor of various patents that are in progress of being issued and to date, he is the inventor or co-inventor of various issued US patents, including:

- 9605529: Prescriptive Reservoir Asset Management - Issued Mar 28, 2017 United States;
- 9230211: Analytics Scripting Systems and Methods Continuation - Issued Jan 5, 2016 United States; and
- 9031889: Analytics Scripting Systems and Methods - Issued May 5, 2015 United States.

Jayendra Pathak

Jay brings a wealth of expertise in malware, phishing, and exploit analysis. He was a research assistant at the University of Houston where he was pursuing his MS degree. A native of Nepal, Jay worked as a computer engineer for the government of Nepal for 4 years prior to moving to the United States. A true researcher, Jay's hobby is to scan the Internet for threats and to determine how those threats affect users. He has a BE in Computer Engineering from Nepal Engineering College and an MS in Management Information Systems from the University of Houston.

Bruce Watson

Bruce is an expert in algorithms for areas such as big data (network security, intelligence, and bioinformatics), pattern matching, computational linguistics, and compilers. Additional work in predictive analytics, domain-specific languages, and decision analysis. Bruce is a highly sought after consultant in the field of Cyber Warfare. He has 2 PhDs in Computer Science.

Kim Coetzer

Kim is B.Sc. Information Technology (with specialisation in Computer Science) graduate, with more than 22 years experience in the delivery of business value and services within complex organisations, delivering outcomes by applying expert knowledge in Data and Information Management, Leadership as well as wide ranging business expertise. His experience has been with a diverse array of business initiatives with a focus on the Financial Services, Mining, Oil & Gas, Energy and Insurance industries. He is well versed in the modern information technology and business environments and has a broad range of skills including Problem Solving, Leadership, People Management, Strategic Insights, Stakeholder Management and Innovation.