Bringing your Security Eco-System closer to Purity utilizing a Vulnerability Data Refinery

June 2014
Author: Gordon MacKay
EVP/Chief Technology Officer

INTRODUCTION

When I was much younger, during the summer following my freshman year in University, I had the opportunity to work for one of Canada's largest gold mines located in Marathon, Ontario. I worked in a lab as part of the metallurgical engineering team and our goal was continuously improving methods which ultimately lead to achieving higher levels of production and concentrations of gold purity. My job consisted of maintaining measurement instrumentation, as well as taking measurements on gold purity at various points within the purification process. What I found interesting is at any given stage of this process, we used multiple measuring devices from different manufacturers so as to ensure the accuracy in the measurements. In other words, the importance of achieving higher levels of gold purity was such a high priority, the process itself required a high degree of measurement accuracy which one vendor's measuring device alone could not provide. As a result, three or four diverse measuring devices from different manufacturers were used at various points in the purification process so as to ensure the trueness of the measured purity value.

The lessons I learned on the importance of approaching trueness of the measured purity of gold, as well as how to achieve it, are applicable to many different areas of life, including the study of information security risk intelligence. In this paper, I explore some of the key challenges organizations face in measuring their information security risk and what is missing from existing security solutions offered within the market in order to achieve higher levels of risk measurement purity. This paper describes a proposed solution that operates as a "refinery". Similar to the processes observed in the gold mine, this refinery imports multiple sources of data which are distilled to the purest form to fuel a highly intelligent information security program.

CHALLENGES IN DEFENDING THE ENTERPRISE

An enterprise's sensitive information represents a substantial part of its overall value. A compromise of this information, even in part, may result in a significant loss in the value, possibly resulting in the enterprise's downfall. Yet, it appears data breach stories are becoming more and more common in the news and media. In fact, according to Symantec's 2014 Internet Security Threat Report[1], 2013 experienced a 62% increase in targeted attacks over 2012. Alongside this increase in attacks, today's enterprises are operating in a more complex environment due to more intricate networks which include BYOD, heightened numbers of mobile devices, an outsourcing of IT infrastructure to cloud services and an increase IT infrastructure access by partners due to an increase in reliance on partnerships. With the above two mentioned forces, organizations are more than ever challenged in defending their valuable information. Unfortunately, organizations also face technological challenges in defending their information. Some of these technological challenges are described next.

Information Value Location Unknown

One of the more impactful challenges enterprises face with regard to security risk is the fact that they are uncertain as to where their valuable information resides and flows within their IT infrastructure. Often, there is no awareness within the organization of the importance of this element of risk, and consequently no attempt to pinpoint the location of their "pure data". As reported, 60% of enterprises are not confident they know where their valuable data resides[2]. As such, when IT professionals analyze the weakness findings identified by the assessment tools utilized, the assumption is all issues must be equal importance. Most often, organizations lack resources to solve all identified "important" issues

and they make little progress in reducing their risk.

Poor Weakness Tracking over Time

One of the most impactful challenges to the accuracy of risk that assessment vendors face is the ability to reconcile assessment findings from one point in time to findings resulting from subsequent similar assessments. For example, a network vulnerability assessment performed January 1st which discovered and assessed 1000 connected network endpoints, must be able to match endpoints discovered in subsequent independent assessments. Assessment vendors are challenged with this because most network vulnerability assessment technology existing within the marketplace today is not agent- based; it detects and discovers the characteristics of the endpoints it is assessing based on the responses it receives from these devices over the network. The characteristics discovered on the endpoints, such as IP address, Hostname, Operating System, MAC address and many more, are all subject to change. To overcome this challenge, most vendors key on one or more of the endpoint characteristics which do not change very often, such as IP address, Hostname or MAC address provided it is detected. However, even this solution is prone to errors and mismatches because IT administrators occasionally move things around in a semi-planned fashion as part of ongoing IT evolution. As a result, the matching is occasionally incorrect. This has several possible dire outcomes such as an endpoint asset being counted twice, or worse, an endpoint asset discovered in one assessment is mismatched to an entirely different asset within a different assessment. I can tell you of "ghost stories" I have heard where IT engineers were chasing the same problem twice, as well as being misled by the technology that reflects many weaknesses for an asset were mysteriously remediated but at the same time, also showed many new weaknesses were discovered. Without exception, all of these stories had one element in common; a very frustrated IT organization along with wasted time and money chasing ghosts.

Imperfect Measuring Devices

Though enterprises generally employ some automated assessment technology so as to understand the weaknesses that exist within the IT infrastructure which holds and protects the information assets, it is important to realize that any given single assessment technology does not provide full assessment coverage. Even worse, for what they do cover, these technologies will mistakenly miss weaknesses that are actually present (false negatives) and will also mistakenly conclude a weakness is present when it is actually not (false positive). In essence, these technologies measure the weaknesses but the measurements may be polluted with erroneous instances, thereby rendering the findings suspect and lacking in purity. This situation illustrates the challenges described in measuring the levels of purity when the measuring devices themselves are imperfect.

No Human Weakness Measurements

As humans, we play a big role within the enterprises we work for and contribute towards. In many senses, a human may actually be viewed as a kind of a computer; an Organic Computer. We hold valuable information such as usernames and passwords. We also operate, use and access the company's computers in order to do work to the benefit of the company we work for. In using and accessing these computers, we may be fooled by others (e.g. in the form of an email) into using a computer to access information that hides a dangerous agent such as malware. We then use this same computer to access valuable Enterprise information, unaware the computer we are using has been

infected. The result could lead to a compromise of the valuable information.

Over time, there has been an increase in cyber-attacks where a human asset was targeted in some form, perhaps as illustrated above, in order to achieve the attacker's goal. In fact, per the 2013 Verizon Data Breach Investigations Report[3], social engineering accounted for 29% of the referenced security breaches. This represents a 400% increase from the previous year. From the attacker's perspective, it is technically easy to achieve and extremely fruitful. Therefore, it is not a surprise that a great deal of enterprise risk is attributable to human weaknesses. Most companies employ an automated vulnerability management technology so as to assess ongoing IT infrastructure weaknesses and prioritize associated remediation. However, we are challenged with the same process as it relates to Organic Computers even though these often hold and access valuable company information.

Technology in a Vacuum

One of the challenges companies face with the existing information security technologies is that they often operate in their own world, within a technology silo. For example, IPS technologies typically apply static rules such as dropping information packets that are deemed as an attack even though the systems they are protecting may not be vulnerable to all of these detected attacks. Ideally in this example, an IPS would use information available from some intelligent system having knowledge of existing weaknesses in order to use this information -to enable the IPS to make more efficient and intelligent decisions in response to an attack.

Consequences of these Challenges

The above challenges ultimately manifest themselves in very costly ways.

Because the whereabouts of valuable information within the organization is uncertain, organizations find it difficult to prioritize remediation of a given IT system over another. As a result, they are overwhelmed with the task of remediating all serious weaknesses for all systems. Since this is costly, choices are made and often they are incorrect, possibly leaving a valuable information asset very vulnerable.

One of the more serious consequences of the impurities resulting from the use of imperfect assessment technology is the compounding effect of assessment network endpoint mismatches. As time progresses, more and more duplicate endpoints "exist" within the IT infrastructure and others are mismatched resulting in weakness turnover nightmare. The measuring devices are so far off that time is required to fully recalibrate. Within vulnerability management, this would be equivalent to abandoning all historical assessment data in favor of a complete new set of baseline assessments.

With the challenge of weakness false positives and false negatives resulting from the use of a single imperfect assessment technology, organizations waste money and yet are missing weaknesses they should have otherwise uncovered.

An organization may follow a solid mature security risk management process for their IT infrastructure but they forget about the human element. Unfortunately, even if the IT arm of the company is security bullet proof, on the infrastructure, attackers will target the employees for credentials. Attackers will also phish employees which could result in a compromise through the installation of malware onto the

employee workstation or laptop. Both cases may lead to a compromise on high information value.

All of the previous mentioned technology challenges result in an organization's uncertainty in security risk and highlight the need for strong security risk intelligence that permeates the entire information security defense strategy.

INFORMATION SECURITY RISK INTELLIGENCE

If we look more closely at the above described challenges, we see that these are all related to a lacking use or availability of security risk intelligence, as well as an impurity of the measurement of security risk. It's impure because though it is considered, it is missing information and for the information that is present, it is erroneous due to the reliance on a limited number of assessment technologies and simplistic ongoing assessment endpoint matching technology.

What is Risk?

> **72%** of Ponemon Institute research[4] respondents identify the **inability to prioritize threats and vulnerabilities** as a challenge that keeps their organizations from being fully effective.
>
> *Ponemon Institute Focus Group Analysis on the VDR Conceptual Design: Preliminary Results*

What is security risk intelligence and why is it important? Understanding and measuring security risk is complex. At a high level, it includes three variables which are interdependent upon each other even though often we treat them as independent. These variables are Value, Threat and Weakness. Value represents the worth of the information the organization needs and generates in order to operate such as employee records, customer information and so on. If there is no value, there is no risk of loss. Threat represents the events that endanger and could possibly compromise the value. Threats may or may not involve an actual human actor with malicious intent. Weakness represents any vulnerabilities that exist within the organization that surround and protect the information from the potential threats. When Value, Threat and Weakness intersect, risk is present. Security risk intelligence is a term which implies information security risk is understood, measured and available for use when needed.

Ideal Security Risk Intelligence System

At a high level, wouldn't it be wonderful to have a powerful entity which is always available and which could answer any security risk question it was posed? I imagine a system where any authorized human and third party technology, could receive a true and accurate answer to question scenarios such as "What is the probability of stored credit card information being stolen in a situation where Bob logs onto system X at 4 am from China while he is using his iPad?" or "What is the probability of my employee data being changed due to an SQL Injection attack against my internal employee website?", or many other questions.

In the first scenario, it is possible Bob is really someone else with Bob's credentials who has mal-intent. Alternatively, it is possible Bob is truly Bob and that he has no malicious intent, but his iPad has been compromised and when Bob accesses the data, it may be passed along to someone else, unbeknownst to Bob.

In the second - scenario, it is possible the internal employee website is not vulnerable to an SQL-Injection attack and therefore, the IPS need not take a protective action which could result in some loss of service.

Ideally, security risk intelligence would be available throughout the enterprise which could provide answers to such risk related question scenarios. With respect to the technology in a vacuum challenge previously described, it would be wonderful if third party technology could reach out to such a dream system, ask it security risk related questions and take appropriate action based on the answer. The following figure 1 illustrates an information security technology product's operation without the help of a Security Risk Intelligence System.

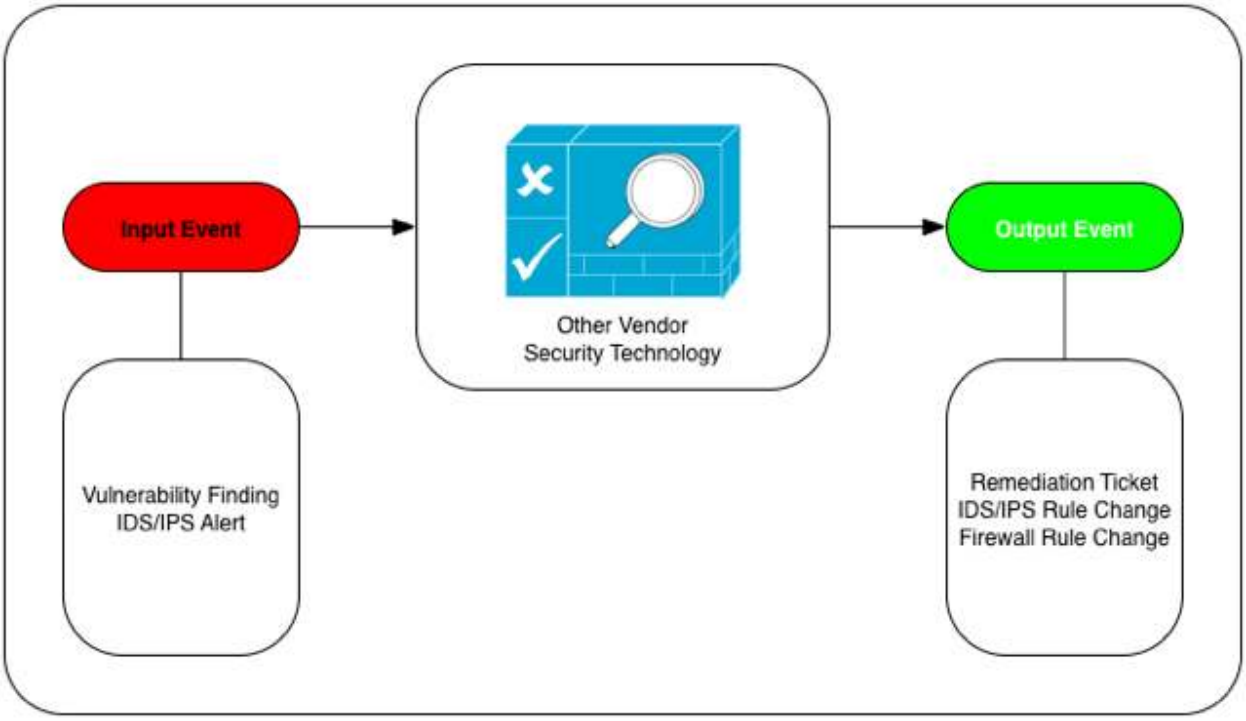

Figure 1 Security Technology Operating in Silo

Figure 2 illustrates the same information security technology in operation, but where it reaches out to a Security Risk Intelligence System that operates as a Vulnerability Data Refinery (VDR) to query for information which helps the security product operate more efficiently. With VDR technology the 3rd party security technology product operates at far greater levels of effectiveness as reflected by the gold output below.

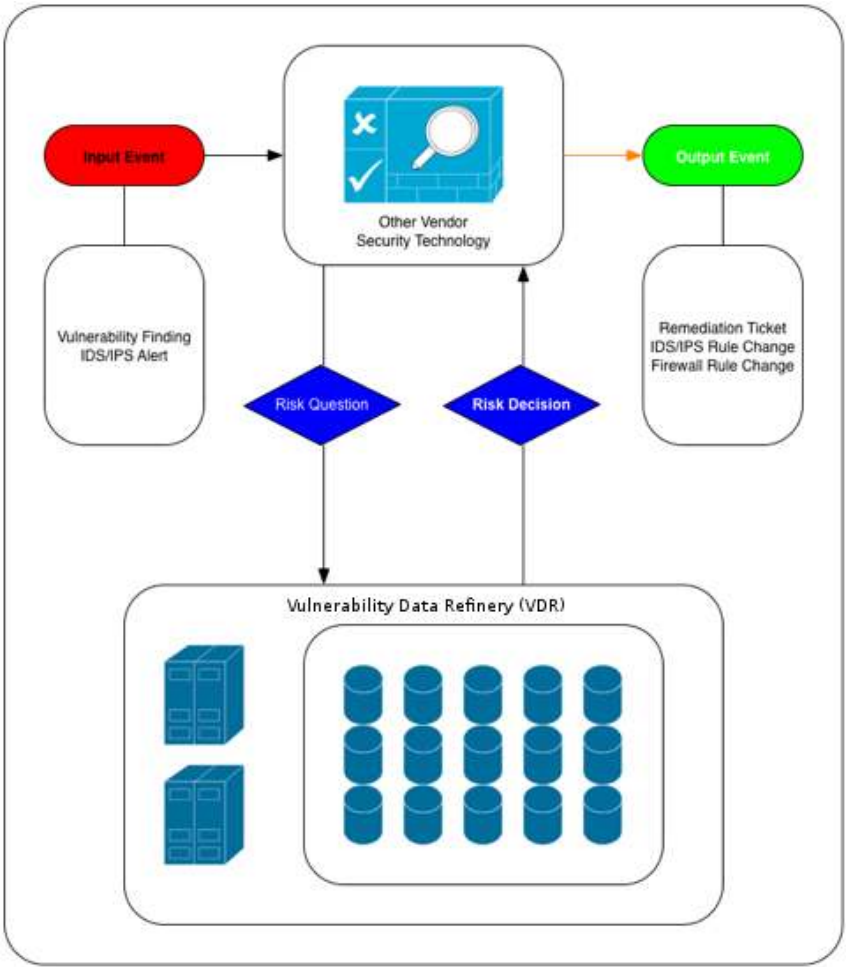

Figure 2 Security Technology Operating with Intelligence from VDR

Revisiting the second question scenario described above, an IPS system that detects an SQL Injection attack could then query a Security Risk Intelligence System and ask whether or not the threat presents any risk, and if so, how much. Integrating with the VDR allows an information security technology product to operate more intelligently.

These above situations highlight that the ideal system must have an interface so that humans and machines can ask these questions and receive answers. It must also have the ability to know about an organization's information value and where it resides. It must have an up to date and accurate view on the weaknesses of the infrastructure as well as for human weaknesses. It must also include consideration for the third variable; security threats.

Knowledge of Risk Triad – Value, Weakness and Threat

As previously described, security risk involves an information value component. Within the context of today's information security landscape, the location of this value within an organization is typically relatively stable over time. Once the location of the value of the information has been identified, it is very unlikely to change from moment to moment. Human and IT weaknesses are also relatively unchanging over time. True, weaknesses are remediated and awareness of new weaknesses emerges over time, but this does not change every second. However, risk does include an element which changes at every moment; this is the context of the threat. We see in the above first scenario question example that Bob is logging on from China at a specific time of the day. The location of the login and the time of day of login are examples of context. This is context which may influence risk and which is directly related to the unique point in time login event. I highlight this distinction because if our ideal system is truly able to answer any risk question posed, it must not only have access to variables which change relatively slowly over time, it must also have access to the information that influences risk and which is only available at the point in time when the given event occurs. With this understanding, the ideal system must somehow have access to contextual threat event information. There are various ways to design such an ideal system such that it includes this capability. I don't explore any of these design options within this section because I want to leave this section open and allow space for the reader's creativity in exploring these.

> VDR's mathematical to risk relies on three components: Weaknesses, Threats from intelligence feeds and Value associated with particular information assets. **90%** of Ponemon research[4] respondents **rated these attributes as IMPORTANT** (59% *very important*; 31% *important*).

Knowing Your Value

In order to gauge relative information risk related to the various parts of an organization, the ideal system must know where and/or with whom the information lies within the organization as well as the relative importance of this value. The ideal system must therefore either allow importation of this knowledge from some external system, and/or must provide the ability for users to specify this knowledge, and/or must provide a capability to discover this knowledge.

Purified Weakness Knowledge

The ideal system must include a very intelligent and self-correcting assessment endpoint matching capability in order to avoid the nightmares related to what I previously referred to as "chasing ghosts." One way to achieve this is where many more of the endpoint characteristics are used within the matching intelligence, much like fingerprint matching algorithms do.

Also, in order to achieve higher levels of purity, the ideal system must overcome the challenge of the imperfect assessment "measuring devices." One way to achieve this is to include a capability where

multiple assessment technologies are sourced and all findings are then correlated together.

Ideal System, Vulnerability Data Refinery, Includes Organic Computer Weaknesses

It is time for us to recognize and admit that a large component of information container weakness is accountable to human users. An advanced Vulnerability Data Refinery should be able to quantify the susceptibility of a given human relative to a given type social engineering attack. The study of how human behavior impacts Enterprise risk is quite young at this time. Still, much research is ongoing and we do have existing qualitative studies which provide us with insight into how this may be assessed. Given the rise in the related threat probability we have seen this past year, the time has come where the human element must be considered.

100% of Ponemon research[4] respondents rated VDR's capability to import data from various vulnerability scanning solutions, application assessments, threat data feeds and human factor risk assessments as **IMPORTANT** (72% *very important*; 28% *important*).

Threat Knowledge – Past, Present and Future

As previously mentioned, one of the variables of security risk is the Threat. Threat is difficult to clearly define on its own without fully understanding the other elements of Risk; Value and Weakness. In the context of security risk, a threat is the probability of occurrence of a "specific" event or scenario. A threat is not any and all events but only those events which have the potential to cause loss of value. A threat may or may not involve a sentient being with specific intent as its source. Typically within the security industry, we discuss threats either in the context of reactive actions or proactive planning.

Protective security technology often reacts and takes action to threat events which occur in the present moment. For example, when an IPS device detects an incoming SQL Injection attack, that event is happening at the given present moment and its probability of occurring is 100% since it is occurring at this moment. At first glance, a simple security risk intelligence system need not know about threats that manifest in the present moment, such as the SQL-Injection threat example just mentioned, because the action of neutralizing such a threat is the responsibility of another device (IPS in this case) and that device already knows about the present moment threat and doesn't need anything else to inform it of this. Following this, when a security protection product such as an IPS asks a question such as "What is the probability of exploitation of device 'X' when subjected to the threat of SQL-Injection attack?", it really only needs to know whether the device under attack is vulnerable to SQL-Injection.

However, examining this a little closer, a more advanced Security Risk Intelligence System, the VDR, would want to know about present moment threats because this would influence its view on the probability of the same threat occurring in the near future. In other words, given part of a network is currently experiencing an SQL-Injection attack, our human intuitive intelligence leads us to assume it is likely the attacks would continue and would possibly target other parts of the network. With this view, the VDR would immediately adjust its view of the threat probabilities related to other devices, in response to current moment threats. It could then issue alerts to humans and/or integrated third party security technology, assuming the system includes risk threshold settings. In essence, VDR would either

have the ability to predict future threat events if it were aware of present moment threats, or it would source this intelligence from an external threat intelligence system.

VDR USE CASES

The VDR described above is of great value to many different types of information security technologies. I provide simple example of how an IPS system could dynamically accept or drop packets based on the qualitative risk responses it receives from VDR to risk questions it poses. Next, I describe two more use cases to illustrate the value of VDR.

VDR Enhances IAM

The following figure illustrates how VDR enhances an Identity Access Management (IAM) System.

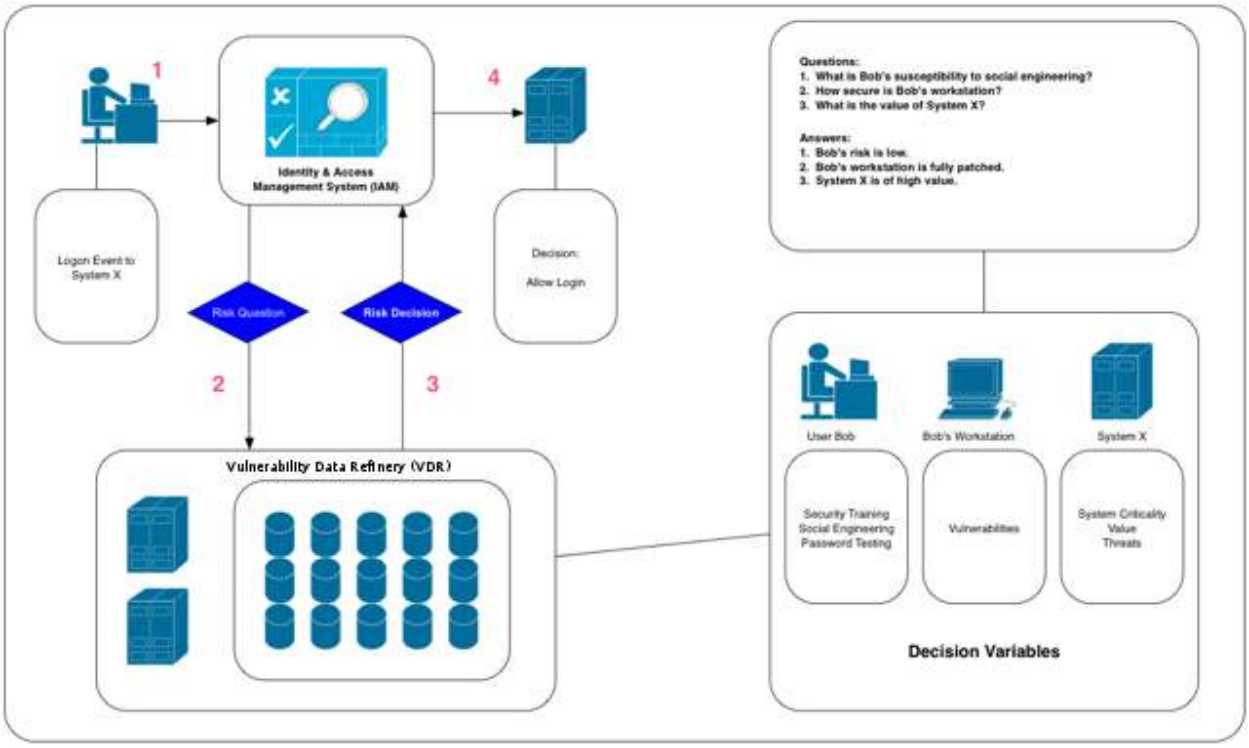

Figure 3 VDR Enhances IAM

This use case assumes that the VDR has retrieved the identities managed by the organization's IAM and that these have been mapped by some mechanism, to the organization's users who have been assessed for security awareness weaknesses.

The use case works as follows:

1- When a user attempts to login to one of the assets operating within the Enterprise, the login attempt is handled by the IAM system.
2- The IAM system would then request the Security Awareness Level for the given user from the VDR. In addition, the IAM system may request risk information associated with the given asset the user is trying to access as well as for the endpoint the user is using to login.
3- VDR would return the information to the requesting IAM.
4- IAM would then know about the risk related to the IT assets as well as the security awareness risk related to the user, and it may use this valuable information during its decision on how to proceed.

Note this is a simplified IAM use case. A more advanced case would include context of the login event which the IAM system would know at the point in time of the login. For example, the IAM would know the IP address of the machine the user is using to login. It would know the time of the login. This contextual information could be passed to VDR as part of the questions posed and it would use this as part of its risk calculations.

VDR Enhances SIEM

The following figure illustrates how the VDR enhances a Security Information and Event Management System - SIEM.

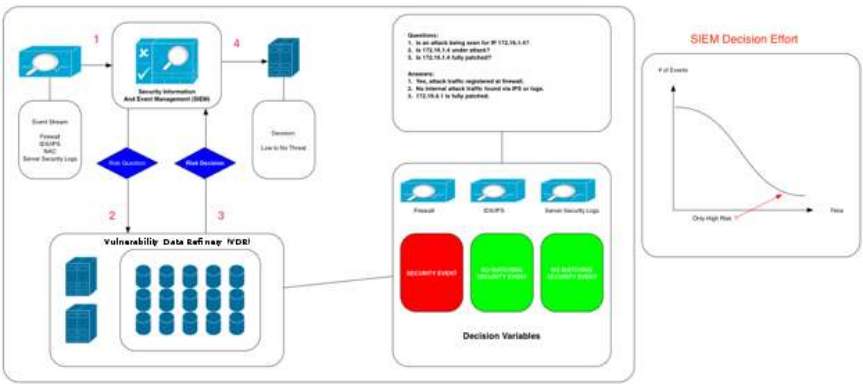

Figure 4 VDR Enhances SIEM

This use case is useful for example by a SIEM offering which employs “Big Data” for security events and where the SIEM product may query the VDR for High Risk Assets and/or High Risk identities. The VDR would then return the asset and/or identity information (including the risk score) that matches the provided risk characteristics. The SIEM product could then filter within all of its events for the given assets and/or identities and thereby allow more focus on these. This is shown in the graph on the right of the above figure where we see the drastic reduction in the number of “interesting” events.

69% of Ponemon research[4] respondents rated VDR's capability to integrate with their organization's existing SIEM and IPS technologies as **IMPORTANT (**41% *very important*; 28% *important*).

CONCLUSION

Security risk intelligence which permeates throughout an organization and which is used by humans and machines enables the organization to operate more efficiently. An ideal system responsible for this intelligence would provide a solution to the impurities resulting from any single existing weakness assessment technology. The advanced solution may achieve this by importing one or more different vulnerability assessment result feeds and by reconciling these results together thereby achieving greater levels of IT infrastructure and application assessment coverage as well as additional confidence in these findings. Optimally, the system must include the human element and must gauge human security weaknesses. To achieve this, the system may have a capability to import results from Learning Management Systems as well as from Social Engineering behavior training systems. Ideally, the solution would be in a position to predicting future risk. This may be accomplished with a capability to import security intelligent threat feeds. Finally, an advanced system must provide an intelligent interface which allows many information security technology products to integrate and retrieve security risk intelligence information thereby greatly enhancing their product offerings.

References

[1]*http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf*

[2]*http://www.varonis.com/assets/reports/en/The-State-of-Data-Protection-Research-Report.pdf*

[3] *2013 Verizon Data Breach Investigations Report ©2013 Verizon*

[4]*Ponemon Institute Focus Group Analysis on the VDR Conceptual Design: Preliminary Results. Commissioned by Digital Defense, Inc. Independently conducted by Ponemon Institute, June 2014*

VULNERABILITY DATA REFINERY (VDR)

USER INTERFACE VIEWS

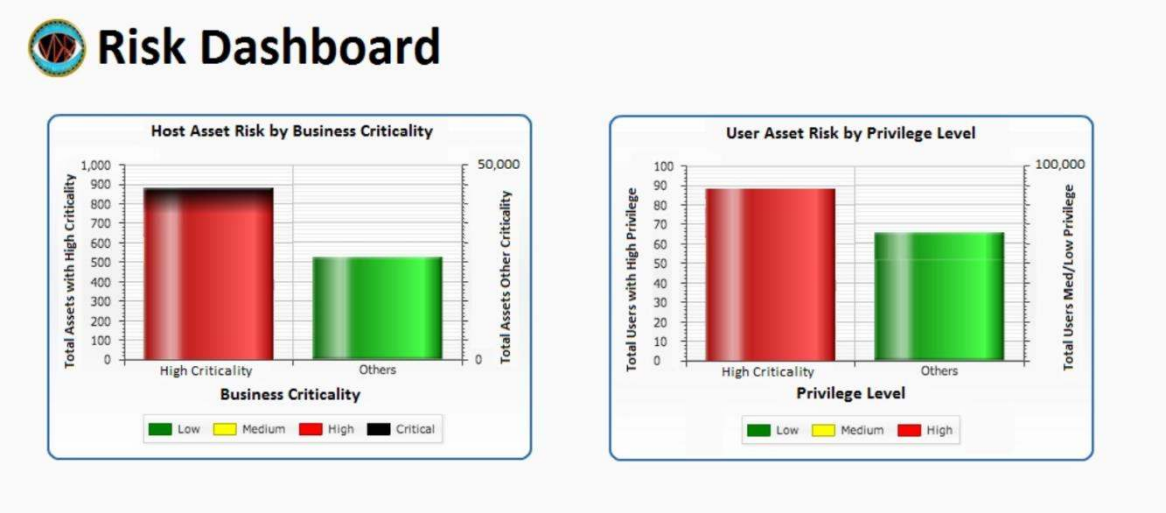

Recent Assessments

Name	Source	Type	Time	Status	Hosts
data network scan	3rd Party	IVA	Jun 12, 2013	completed	6 / 6
London Database Servers		IVA	Jun 13, 2013	completed	11 / 11
Tomcat Easily Guessable Password Rescan		Auth Scan	Jun 13, 2013	completed	7 / 7
Tomcat Easily Guessable Password Rescan2		IVA	Jul 30, 2013	completed	1 / 1
London POS Website DAST	VERACODE	DAST	Jul 30, 2013	completed	1 / 1
London POS Website SAST	VERACODE	SAST	Jul 30, 2013	completed	1 / 1
Frankfurt Firewall Disabled IPSec		IVA	Jul 30, 2013	completed	1 / 1
Clusterix DB Default Password		IVA	Jul 31, 2013	completed	1 / 1
Conficker Fixed		IVA	Jul 31, 2013	completed	1 / 1
Conficker Rescan		IVA	Jul 31, 2013	completed	2 / 2

Show more results

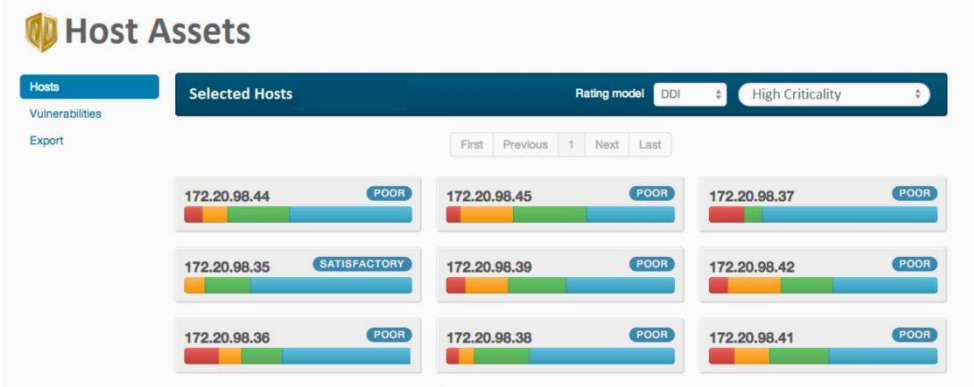
Host Assets
Hosts
Vulnerabilities
Export
Selected Hosts
Rating model
DDI
High Criticality
First
Previous
1
Next
Last
172.20.98.44
POOR
172.20.98.45
POOR
172.20.98.37
POOR
172.20.98.35
SATISFACTORY
172.20.98.39
POOR
172.20.98.42
POOR
172.20.98.36
POOR
172.20.98.38
POOR
172.20.98.41
POOR

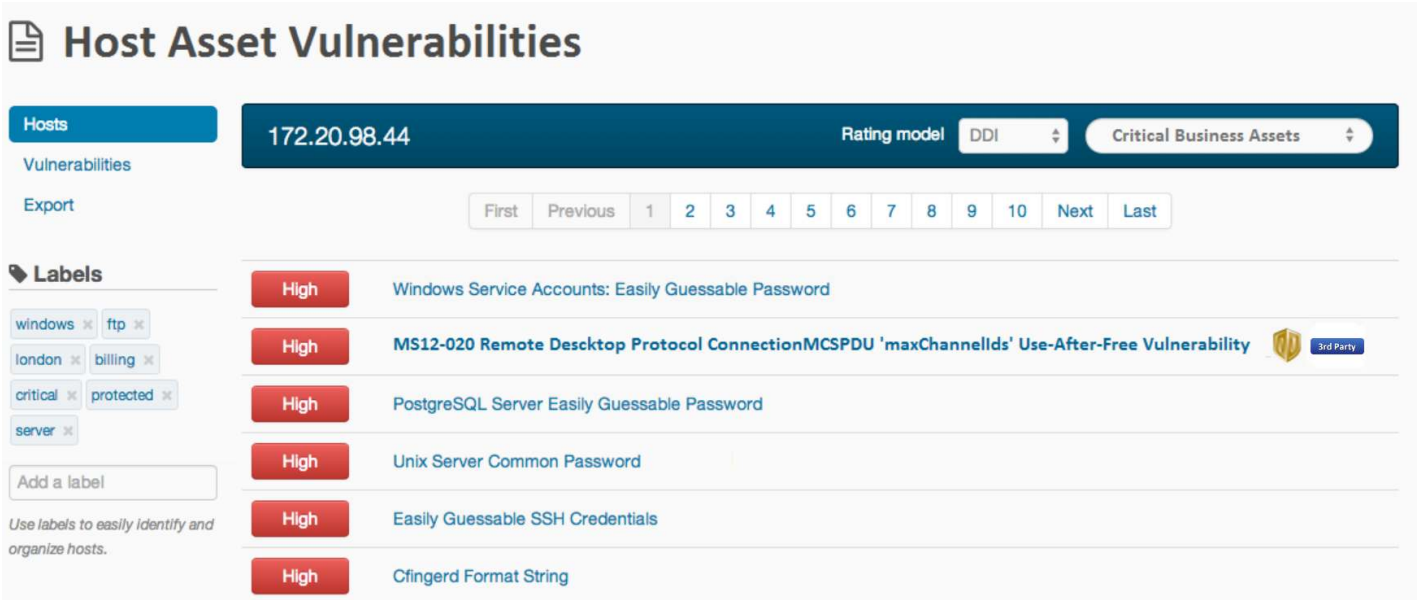
Host Asset Vulnerabilities
Hosts
Vulnerabilities
Export
172.20.98.44
Rating model
DDI
Critical Business Assets
First
Previous
1
2
3
4
5
6
7
8
9
10
Next
Last
Labels
windows
ftp
london
billing
critical
protected
server
Add a label
Use labels to easily identify and organize hosts.
High
Windows Service Accounts: Easily Guessable Password
High
MS12-020 Remote Descktop Protocol ConnectionMCSPDU 'maxChannelIds' Use-After-Free Vulnerability
3rd Party
High
PostgreSQL Server Easily Guessable Password
High
Unix Server Common Password
High
Easily Guessable SSH Credentials
High
Cfingerd Format String

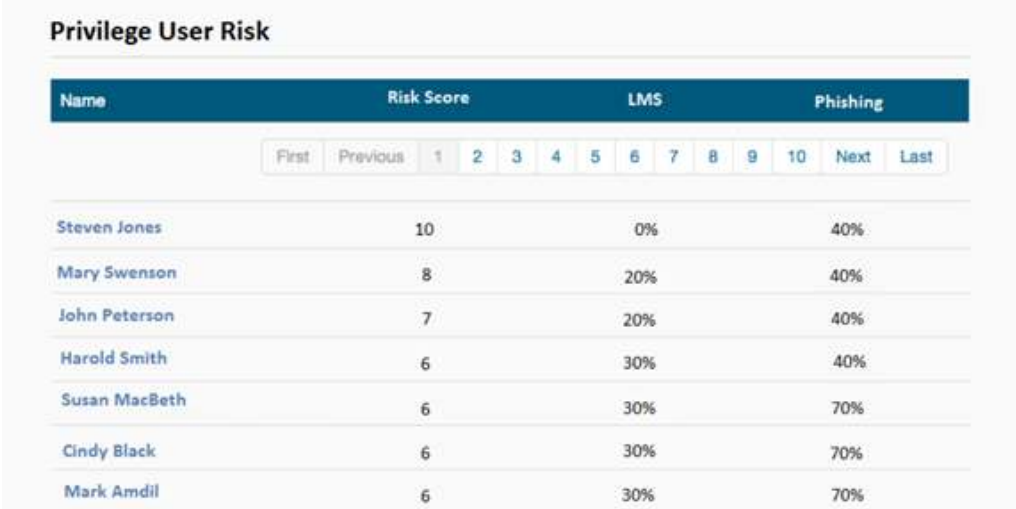
Privilege User Risk
Name
Risk Score
LMS
Phishing
First
Previous
1
2
3
4
5
6
7
8
9
10
Next
Last
Steven Jones
10
0%
40%
Mary Swenson
8
20%
40%
John Peterson
7
20%
40%
Harold Smith
6
30%
40%
Susan MacBeth
6
30%
70%
Cindy Black
6
30%
70%
Mark Amdil
6
30%
70%