A Framework for Business Aligned Security Monitoring Use Cases

Get relevant and improved visibility of your security risks

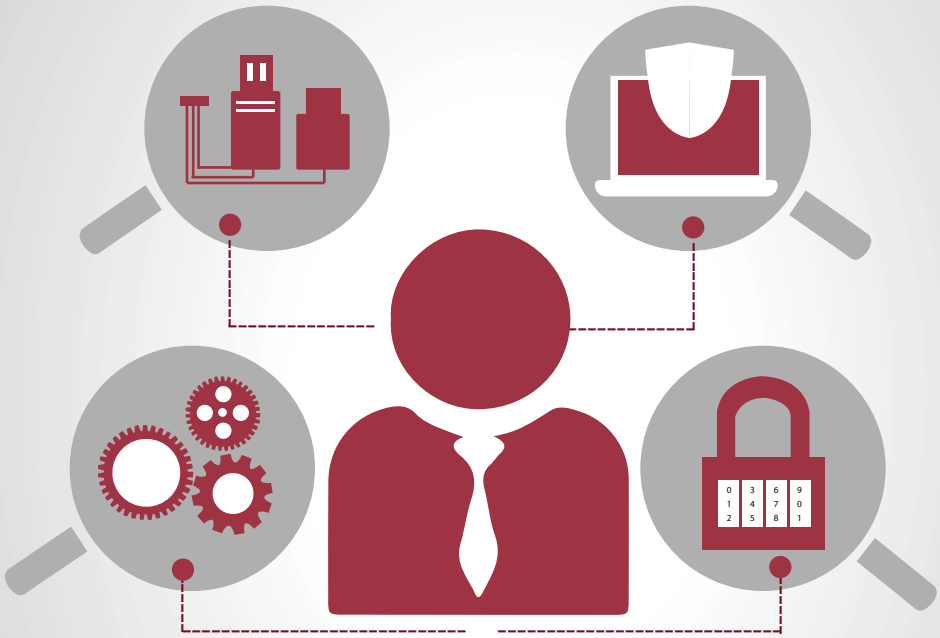

Author:

Vinod Vasudevan, CTO

Contributor:

Sujay Mendon, Research Lead – SOC

PALADION
HIGH SPEED CYBER DEFENSE

Table of Contents

Introduction

"Most organizations prefer to rely on default monitoring rules that are provided by a SIEM vendor. While this approach is a good start, it does not provide the desired level of assurance that the system is monitoring the specific risks that the organization should be watching for."

This white paper makes the case and provides a framework for applying business logic to Security Incident and Event Monitoring (SIEM), thereby improving detection capabilities, focusing resources on the highest impact areas, and demonstrating the business value of security monitoring and operations.

Security monitoring today faces a dichotomy - security teams want to increase the breadth and depth of monitoring to detect threats before there is a breach, while the business managers question the value of increasing the budget.

To a security professional, the justification for enhanced monitoring is evident - attacks and attackers have both evolved in volume and complexity. The solution lies in bringing more assets and more data sources into monitoring, using SIEM and big data analytics together and implementing more use cases. Although it may not be as obvious to the business decision makers, there are enough high profile breaches to make a case.

True security monitoring involves many complex tasks, from collecting logs, constructing a SIEM, developing a security analytics platform to implementing continuous operations for alerting and responding. However, that is not where true value is generated. The common business rule of thumb, the 80/20 rule applies here: 80 percent of the value in security monitoring is generated from the quality of use cases, which account for 20 percent of the security monitoring activities.

Most organizations prefer to rely on default monitoring rules that are provided by a SIEM vendor. While this approach is a good start, it does not provide the desired level of assurance that the system is monitoring the specific risks the organization should be watching for. The question is how use case modeling can be business focused so that the value creation is visible to the management.

We have reviewed over 100 security operation centers, in order to glean best practices for use case development. In the majority of cases, we've seen that the approach to use cases is to apply all readily available rules of SIEM to all log sources. In some cases, new rules have been built, but the approach has been to leverage the SIEM features to their maximum potential rather than focus on what the business needs to protect.

Problems with Current Methodology

This approach inherently leads to problems in both mitigating risks and articulating the value of monitoring for the decision makers:

The alerts that are generated based on these use cases become technical jargon rather than describing the impact or how a threat can materialize in the organization's context. Questions like how critical an asset is to the business process, what the threat will do to the asset, or what the impact will be to the business process from this threat remain unanswered for the business.

The use cases are not derived from the risks the business is potentially exposed to. While technical use cases can be the same for an entire organization, risks are more specific and are based on an organization's context for assets, users, data type, and threat perception. Consequently, many of the risks specific to an organization may not get modeled into the use cases when only the technical ready content is applied.

To protect against a risk, multiple use cases may need to work together and go beyond just the SIEM to other devices like Intrusion Prevention System (IPS), Web Application Firewall (WAF), or big data analytics. A narrow application of a use case to just the SIEM leaves a considerable gap in the detection of the risk even though a technical use case may already exist.

Can there be a different approach to monitoring use cases that can increase technical detection while improving the value to a business? Based on our survey of some of the most advanced security operation centers, we've created a framework with the potential to do so.

"While technical use cases can be the same for an entire organization, risks are more specific and are based on an organization's context for assets, users, data type and threat perception."

Implementing a Risk Based Framework

The framework for building monitoring use cases has two distinct phases: risk modeling and technical rule building.

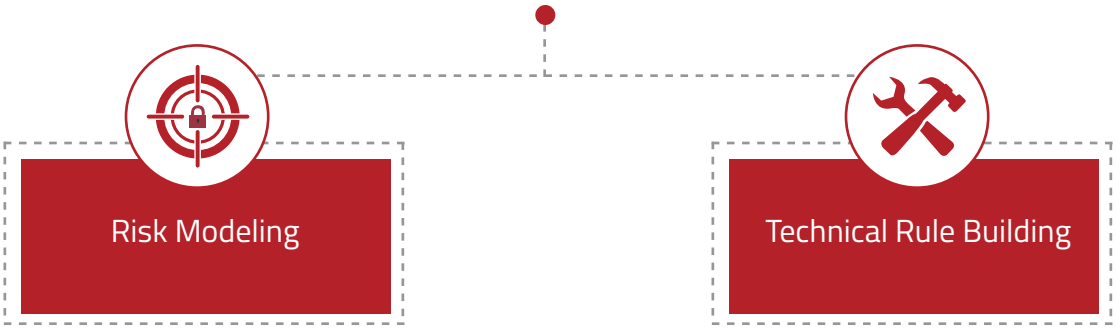

This framework will be able to:

- Make the use cases relevant as well as correspond to the specific risks that businesses are exposed to.
- Apply use cases beyond SIEM. Other security and network devices will also be leveraged to create a well-orchestrated monitoring environment where a risk will trigger alerts at multiple stages for deeper detection.
- Provide greater context to alerts. The use cases will not just generate technical alerts but also contextual risk data that helps to better understand its impact on the business.
- Make the business feel involved in security monitoring and highlight the significance of a SOC to business management.

Risk Modeling

"The first step in risk modeling is to identify the high value business processes or business units that need to be protected."

A business continuity plan is based on a Business Impact Analysis (BIA) that quantifies impact and guides investment decisions. Similarly, setting up an information security management system starts with a risk assessment (RA) that quantifies the risks and guides control selection. In the same manner, the technical rules for monitoring should be based on risk modeling. In our framework, the risk modeling begins by first identifying the business processes, the business risks and the assets to be protected and then creates a repository of risk statements for monitoring.

1. Identify Business Processes

The first step in risk modeling is to identify the high value business processes or business units that need to be protected. If an organization has already conducted a BIA or RA, processes with a high rating are identified. For the purpose of understanding the framework we will consider one business process each in retail and banking. In retail, we will look at a customer retention program and in banking, at an Internet Banking Service.

2. Identify and Record Business Risks

In the next step we record the key business risks for each process. For a retail customer retention program, one of the key risks to consider is customer data theft. For an internet banking service, one of the key risks to consider is channel downtime for customers.

3. Focus on Target Assets

Once the list of business risks is created, we can identify the assets that might be targets for realizing these risks.

Let us apply this in our retail example. Assets that contain customer data need to be the focus since we are looking at a customer retention program with customer data theft as a risk. Hence retail Customer Relationship Management, Enterprise Resource Planning, and Central Retail Application become the target assets for the creation of a technical use case.

For an internet banking service with the risk of channel downtime, the target assets include an internet banking link, router, web server, application server, and database.

4. Identify Threat Actors

The term "threat actor" denotes potential attackers or malicious users who can execute attacks on target assets supporting those processes identified with the highest business risk.

In our retail example, the threat actors for customer data theft are:

- External
  - Third-party service providers
  - Cyber-crime syndicates
  - Competitors
- Internal
  - Administrators
  - Staff
    - Resigned employees serving notice
    - Staff working with CRM/ERP/Retail Apps
  - Contractor/Consultant

In our banking example, external threat actors are more likely to be the cause of an internet banking service outage.

- External
  - Cyber-crime syndicates
  - Competitors
  - Nation States
  - Hacktivists

In order to identify threat actors, a repository of threat actors and vectors can be pre-built to assist in risk scenario planning. Paladion provides a repository in its use case modeling tool, which is provided as freeware to assist in this step.

5. Analyze Threat Vectors

In this stage, we look at the potential threats or attacks that can be executed by threat actors for the estimated business risk.

In our retail example, the threat vectors for customer data theft can be:

- Account Compromise / Takeover
  - Password cracking/Session hijacking
  - User Privilege Manipulation
    - Adding users to administrator group, privilege escalation
    - Role and permission tampering
- Malware Infection
  - Rootkits/ Spyware/ Backdoors/ RAT/ Malicious Code

- Social Engineering
  - Phishing/Spearphishing/ Vishing
  - Watering hole attacks
- Web Based Attacks
  - XSS/SQL injection/CSRF

In the banking example, the threat vectors for internet banking service downtime can be the following:

- DOS
  - DDOS/ DNS reflection/Smurf
  - Host level DOS
- Web Based Attacks
  - XSS/SQL injection/CSRF
- Network Attacks
  - Port scanning, Vulnerability scanning, buffer overflow exploit

6. Articulate Risk Scenario

The final step in risk modeling is to capture the risk scenario using a risk statement. There can be multiple risk statements based on the threat vectors, actors, and business risks.

Sample risk scenario for the retail example:

Business Process	Business Risk	Target Assets	Threat Actors	Threat Vectors
Customer Retention Program	Customer Data Theft	CRM, ERP, Central Retail Application	Competitor	Social Engineering Malware

Risk Scenario:
A competitor uses social engineering and malware to steal customer information from a CRM system leading to customer churn and revenue loss.

Sample risk scenario for the banking example:

Business Process	Business Risk	Target Assets	Threat Actors	Threat Vectors
Internet Banking Service	Internet Banking Channel Down-time	Internet banking link, router, web server, application server and database	Hacktivists	DDOS/DNS Reflection

Risk Scenario:
A hacktivist group launches a DDOS attack to target an internet banking system, making the server unavailable for legitimate users resulting in customer dissatisfaction and reputation damage for the bank.

Creating risk scenarios can be automated to a large extent using risk modeling tools. A freeware tool is also available from Paladion.

Technical Rule Building

The next step for use case modeling is to derive the technical rules for each risk scenario. The method for building technical rules includes four stages:

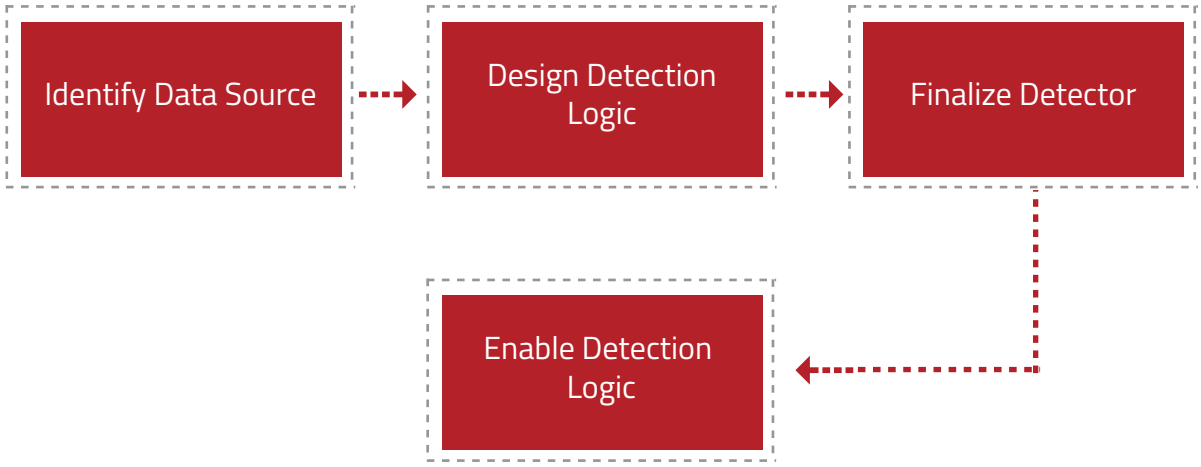

"A good monitoring strategy should utilize every piece of detection ammunition you possess."

1. Identify Data Source

The first stage is to identify the source of data that is relevant for detecting the risk. These sources could be server logs, firewalls, packet capture data, proxy, netflows, user access data, or others. We arrived at the risk statement, "A competitor uses social engineering and malware to steal customer information from a CRM system leading to customer churn and revenue loss" for the retail example that we have been discussing in this paper. To identify sources corresponding to this, we will need to look at the OS, Webservers, and databases corresponding to the CRM system. We will also need to look at logs from firewalls, proxy, and network devices surrounding the CRM system. In addition, there may be more components to be monitored depending on the network architecture and the user access to the CRM system.

2. Design Detection Logic

The next stage is to determine what detection logic should be used to indicate that a risk has occurred.

In this case, this is a logical statement rather than a technical statement. Returning to the retail example, we will need to select detection logic that relates to malware, watering hole (social engineering), and data exfiltration.

3. Finalize Detector

Next, we will need to identify where the detection logic should be applied for best results. A SIEM is not the only product that should be used for detection of risks. A good monitoring strategy should utilize every piece of detection ammunition you possess, which would mean leveraging technologies including IPS, WAF, Data Leak Prevention (DLP), Network Behavior Anomaly Detection, big data analytics, and others. Considering the retail example to determine how we can detect malware and social engineering attempts, the SIEM can be used for known malware and unknown malware can be detected using big data analytics. WAF can also be used as the detector for some of the leading indicators of a malware attack including SQL injection.

4. Enable Detection Logic

Finally, we will need to translate the detection logic into rules or analytical outcomes based on the detector. Each detector will need its own specific technical rules. In current implementations across most SOC's, use case building is only concerned with this aspect, whereas it should be a culmination of the entire risk modeling and technical modeling process.

Each risk scenario will have multiple technical rules. For example, the retail risk scenario can be converted to technical rules as shown below.

Risk Scenario: A competitor uses social engineering and malware to steal customer information from the CRM system leading to customer churn and revenue loss

Source data	Detector	Detection Logic
Web Traffic	WAF	SQL injection attack for malware implant. This is a rule to be enabled in WAF
Proxy	Analytics System	Detect proxy traffic logs with entropy value of near zero which indicates low "uncertainty of data". This indicates systems infected with malware. This is an example of applying analytics for a risk scenario.
Network traffic	DLP	Keyword search for confidential documents. This is a rule configured in DLP.

Making Sense of Correlated Rules

"To build correlated rules in a systematic manner derived from risk, the main question to ask is "What detection logic or conditions will validate the risk that has materialized?"

Rules like the ones established above identify whenever a risk scenario is likely to materialize. These rules can also be triggered by normal operations, sometimes leading to false positives. There are two ways to enhance these technical rules through higher order rules or correlated rules. The security industry has already built several correlated rules in SIEM, however these are not built through a structured method applying the business risk.

To build correlated rules in a systematic manner derived from risk, the main question to ask is "What detection logic or conditions will validate the risk that has materialized?" This question is different from "What detection logic will indicate the risk scenario?" which refers to building technical rules from indicators as per the table above.

Each risk scenario can have one or more validators. Continuing with the same example, the validators and correlated rules are:

Risk Scenario: A competitor uses social engineering and malware to steal customer information from CRM systems leading to customer churn and revenue loss.

Detector	Validation Logic	Correlated Rules
Analytics System	DLP keyword search reveals a set of users with data exfiltration. Validate by checking which of the users show abnormal high outgoing traffic.	DLP alert on users based on keyword match combined with analytics alert for abnormal high outgoing traffic for the users.
SIEM	C&C traffic alert in SIEM for an asset based on threat intelligence feed. Validate to see if the same asset had a successful SQL injection alert	Correlate SQL alert on an asset combined with C&C traffic match from the same asset.

Linking Multiple Risks to Detect Attack Campaigns

The next level of correlation involves linking different risks and determining if there is an attack campaign in progress. This goes beyond building rules for each risk scenario as the aim is to correlate a longer period of risk data to determine if targeted attackers are running a campaign against an organization.

This can be achieved by tagging each risk to one or more of the cyber kill chain stages. Over a period of time the risks for every asset, as well as the alerts corresponding to the risks, can be seen as a progression along the cyber kill chain. These potential kill chains give greater visibility to security analysts and can be analyzed for deeper investigation into a potential attack campaign.

An attack campaign related to the retail risk scenario "A competitor uses social engineering and malware to steal customer information from CRM system leading to customer churn and revenue loss" can be detected as follows:

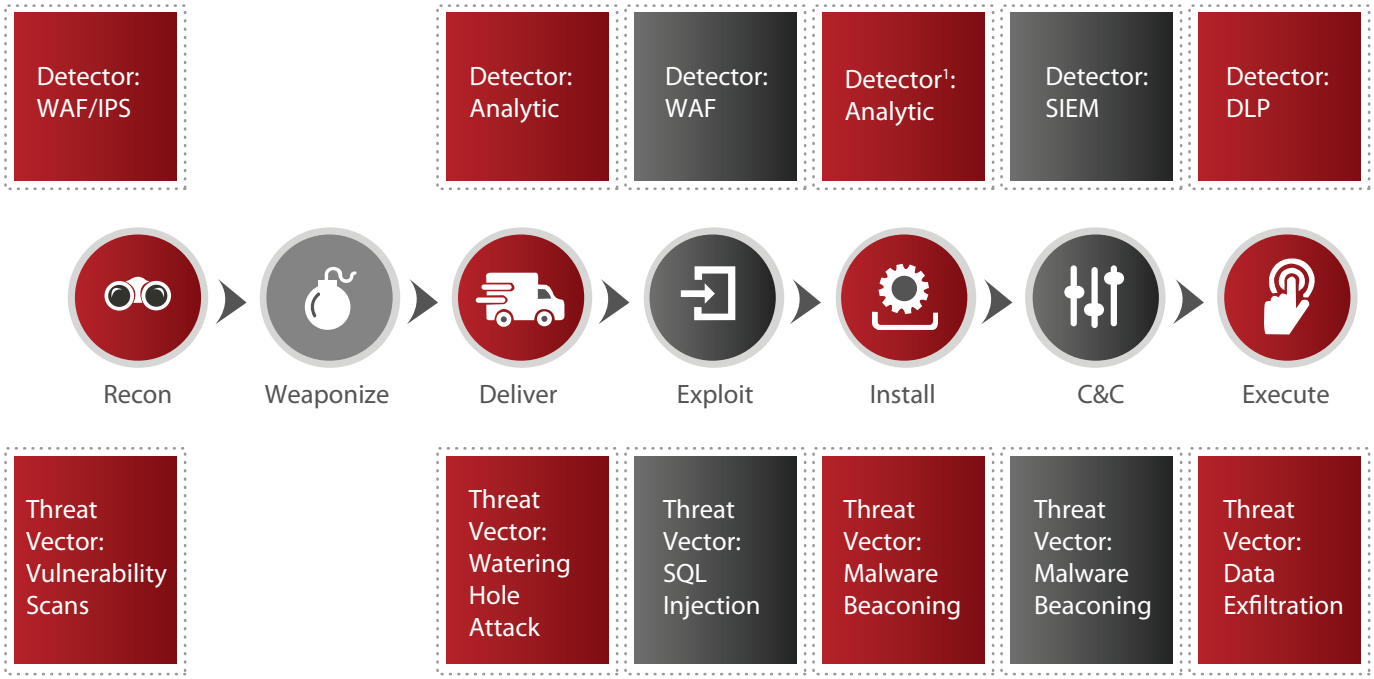

Output of Montoring

Once the rules are built into detectors (SIEM, Analytics, WAF, etc), the alerts generated are converted back into the actual risk scenario. Such conversions can be carried out through a central alert reporting tool which may be accessed by business owners as well as a security team in an organization. The Paladion RisqVu platform provides this capability for risk modeling, technical rule generation as well as reporting alerts in business risk language. It is no longer in the purely technical language of attack names and IP addresses.

Summary

The effectiveness of security monitoring is primarily dependent on the quality of use cases. The quality should consider identifying the risks from a business perspective and building a more comprehensive detection of the identified risks across a variety of security technologies including SIEM. This approach is different from applying generic readily available rules in SIEM for monitoring all logs put into SIEM. The focus of this whitepaper has been to describe a model for achieving high quality business focused use cases for monitoring. The model can be executed with the help of automated tools and a risk repository, which are available as freeware from Paladion.

Glossary

Data Source: Defined as a system or an application component that generates logs and records any process, operation, change or interaction with that system. Network devices such as routers/switches, web servers, operating systems, databases and application servers are some examples.

Detector: Defined as a device or a system that generates alerts (warnings, error notifications or informational messages) in response to a specific activity based on in-built logic and rules. Firewall, IPS, WAF and DAM are examples of detectors.

Impact: Defined as the effect or consequence that results from something such as data loss, data theft, system outage or system unavailability that could directly impact the business, such as revenue loss, customer dissatisfaction, damage to reputation, productivity or efficiency loss, legal and compliance penalties, etc.

Indicators: Defined as logical entities or objects such as virus or intrusion signatures, IP addresses, domain URLs, email IDs, or events, such as login account failures, high CPU utilization, TCP fragmentation, etc, that are associated with or point to malicious or suspicious behavior, including an attack.

Rule: Defined as a logical expression or conditional construct designed to detect events of interest from a security standpoint that warrant an investigation. Rules are implemented on detectors such as IPS or WAF, or on security monitoring systems such as SIEM or log analytics platforms. Example of a rule defined in SIEM: Send an email alert to the application owner if there are login failures for an internet banking account from multiple geographies within 10 minutes.

Threat Actor: Defined as an individual, group, organization or government that conducts or has the intent to conduct detrimental activities. Threat actors could be administrators, contractors, employees, business partners, external attackers, hacktivists, rogue states or nation states.

Threat Target: Defined as a person or an asset. Assets can be further grouped into business units and users can be grouped by departments, locations, etc.

Threat Vector: Defined as a method the threat actor uses to reach a target. Some examples of threat vectors are social engineering, account takeover, data theft, DOS, fraud, malware, network/physical/web attacks, policy violations, system sabotage and system manipulation.

Use Case: Defined as a mechanism to detect and/or confirm a possible security compromise, a breach or a policy violation. A use case can be realized with a combination of rules and/or dashboards and/or reports configured within a detector.

PALADION
HIGH SPEED CYBER DEFENSE

ABOUT PALADION

Paladion is a global cyber defense company that provides Managed Detection and Response Services, DevOps Security, Cyber Forensics, Incident Response, and more by tightly bundling its semi-autonomous cyber platform and managed services with leading security technologies. Paladion is consistently rated and recognized by independent analyst firms and awarded by CRN, Asian Banker, Red Herring, amongst others.

For 17 years, Paladion has been actively managing cyber risk for over 700 customers from its six cyber operations centers placed across the globe. It houses 900+ cyber security professionals including security researchers, threat hunters, ethical hackers, incident responders, solution architects, consultants and more. Paladion is also actively involved in several information security research forums such as OWASP, and has authored several books on security monitoring, application security, and more.

WW Headquarters: 11480 Commerce Park Drive, Suite 210, Reston, VA 20191 USA. Ph: +1-844-507-7668

Bangalore: +91-80-42543444, Mumbai: +91-2233655151, Delhi: +91-9910301180, London: +44(0)2071487475, Dubai: +971-4-2595526, Sharjah: +971-50-8344863, Doha: +97433559018, Riyadh: +966(0)114725163, Muscat: +968 99383575, Kuala Lumpur: +60-3-7660-4988, Bangkok: +66 23093650-51, Jalan Kedoya Raya: +62-8111664399.

sales@paladion.net | www.paladion.net