MANDIANT®

A FireEye™ Company

WHITE PAPER

Using Metrics to Mature Incident Response Capabilities

Compiled by:
Mandiant Security Consulting Services

April 9, 2014

CONTENTS

USING METRICS TO MATURE INCIDENT RESPONSE CAPABILITIES

INTRODUCTION

Incident response requires a symbiotic relationship between people, process, and technology. Organizations can take several recommended steps to diminish the threat; however, the risk of compromise can never be eliminated. Knowing this, Mandiant recommends that companies focus on improving the cycle time for their incident response process, which is defined as a function of the time between initial compromise and threat containment. To this end, metrics play a key role in maturing an organization's incident response capability.

DRAIN CVR

Most simply, Mandiant calculates the Dwell Time and the Containment Time of an incident. The Dwell Time refers to the length of time from the initial compromise through the point of notifying affected stakeholders. The Containment Time refers to the period between collecting live response data and the eventual remediation. Organizations may choose to define remediation as simple containment (preventing the threat from communicating or moving laterally) or full restoration of service (threat eliminated and systems returned to normal).

These two (2) categories are broad in nature and can be further segmented and studied as an organization's incident response capability matures. The Dwell metric can be broken down into the following component parts: Detect, Review, Analyze, Identify, and Notify (DRAIN). Similarly, the Containment portion can be measured using the sub-metrics Collect, Validate, and React (CVR). Each component of Mandiant's DRAIN CVR metric system is defined in Table 1.

Time	Category	Measurement	Benefit
Dwell	Detect	Time from initial entry into the system/network to detection	Measures the effectiveness of detection systems and capabilities
	Review	Time from detection of the incident to analyst for review	Determines if staffing level is properly sized
	Analyze	Time to analyze the incident	Determines if the organization has the right expertise and tools and if the right escalation occurs
	Identify	Time to identify the affected assets, location, and owner	Measures the effectiveness of asset inventory
	Notify	Time to successfully notify appropriate contacts	Measures the effectiveness of contact database and communication plan
Containment	Collect	Time to collect live response data	Determines if the right tools are deployed to assist in collection
	Validate	Time to validate intrusion based on collected data	Determines if the right skill sets are in place at each level
	React	Time to react (contain, remove, etc.)	Determines if the right definition of remediation exists and if it is applied consistently

Table 1: DRAIN CVR Definitions

These metrics successfully measure the efficiency of people, process, and technology using a succinct system that can eventually be automated in most organizations. When an organization reaches a maturity level that allows for full DRAIN CVR measurement, it becomes easy to identify bottlenecks — places where the process stops or slows — within the overall incident response process. Understanding process bottlenecks can influence asset allocation (perhaps toward improving hiring, training, process development, or technologies) in an effort to develop a stronger incident response capability.

ADDITIONAL METRICS

The DRAIN CVR metric is most directly related to detecting, responding to, and containing a threat — the core components of incident response. However, Mandiant's Security Consulting Services understands that the most effective Security Operations Centers (SOCs) and Computer Incident Response Teams (CIRTs) benefit from a tactical and holistic programmatic approach. This means not only calculating the DRAIN CVR metric, but also measuring how effectively the CIRT is implementing innovative solutions as well as interacting with the entire business. The information security community has divided metrics into three main categories: efficiency, implementation, and impact.[1]

Efficiency Metrics

Efficiency metrics help a security organization understand its own performance. For example, the DRAIN CVR approach described above is an approach to structuring efficiency metrics for a CIRT, but there are many other efficiency measurements that a security team might employ.

Each element of DRAIN CVR is itself an efficiency metric. By measuring each of them, an organization can understand its capability in each area and, by breaking that element down even further, can learn how to improve. For example, if the initial detection phase is a bottleneck, an incident response team might ask whether they've properly optimized their network, if they have appropriate indicators of compromise (IOCs) in place, or if their technology is capable of detecting today's threats.

Implementation Metrics

Implementation metrics are less detailed than efficiency metrics, and are often reported to senior leadership. They allow the CIRT to focus on how well new programs and processes are being introduced. While efficiency and implementation metrics may focus on similar incident response processes and security controls, the two are approaching these activities from very different angles.

Implementation metrics are intended to determine the progress of different practices being newly applied to an environment. Ideally, there will be established goals considering these new implementations. This metric will demonstrate whether or not the CIRT is meeting implementation goals and benchmarks. This category can focus on CIRT internal activity, or it may be used to measure new procedures governing CIRT interaction with external stakeholders.

A common implementation metric may measure the rollout of a new host-based security technology to an environment. Several factors, including access to each host and proper configuration of the agent, affect the overall success of the enterprise deployment. An organization can measure the success of the technology deployment using implementation metrics. In this case, the measure begins at zero percent and is completed at one hundred percent. Implementation metrics are unique in that they have a limited lifespan. In this case the host-based technology can still be measured, but after successful rollout, the metric transitions from that of implementation to efficiency. This same process is not limited to technology deployment, but can also be used to gauge the success of integrating new processes and procedures.

[1] NIST SP 800-55 Rev 1. 2008.

Impact Metrics

Finally, impact metrics are relevant for the highest level of governance. These measurements help to illustrate the overall impact that the CIRT has on the business as a whole. Efficiency and implementation metrics are primarily meant to inform CIRT teams, whereas impact metrics consider executive leadership to be their primary audience.

A security organization will demonstrate cost savings, brand protection, compliance achievements, or other business value to leadership using impact metrics. These metrics will often use information gained from efficiency or implementation metrics and portray that data against a set of resources to demonstrate value. Impact metrics may track a trend over more extended periods of time and are not likely to produce valuable data as quickly as efficiency or implementation metrics.

An example of an impact metric might be “Incidents per CIRT person-hour,” which would explain to management the impact of the new technologies they improved in the last budget cycle, or “Incident Cycle Time,” to illustrate the impact of the new full-time analyst.

Proxy Metrics

Proxy metrics can fall into any of the previously described categories and are not a category of measurement on their own. Rather, proxy metrics are easily gathered and allow an organization to measure those seemingly intangible security issues that are often the most important.

For example, many mature CIRT environments may not place a high face value on measuring every Microsoft patch deployment. However, success in this category can demonstrate a high degree of control over the organization’s assets and enterprise network – a quality all CIRT managers desire throughout detection, response, and containment.

In another example, the number of antivirus alerts per asset may seem insignificant, especially because we know that targeted attackers are using malware on only half of the assets they compromise.[2] But this metric, too, provides an organization with a way to understand the overall hygiene of an environment.

[2] *M-Trends 2012: An Evolving Threat*. Mandiant.

COLLECTING METRICS

Determining What to Measure

In general, metrics are not difficult to calculate. Meaningful metrics, however, can be. As with any project it is important that a CIRT interested in maturing their ability to use metrics do so from a logical starting point. By flowing requirements downhill from the highest-level business objectives, an organization can go from a high-level business goal to a focused and quantifiable metric in five (5) simple steps.

Imagine that a financial services company set the goal of being the number one credit card provider in the country. The security implication of this goal would be the need to maintain trust. From this implication, an IT department may develop a policy requiring that all applications adhere to software security standards. Logically continuing this downhill requirement flow, a security standard may read: “All input from web applications must be properly validated per this spec.” From this standard, the CIRT is able to measure the percent of applications found to be vulnerable to SQL injection. In this example, the metric is easily derived from, and logically supports, overall business goals and objectives that otherwise may have appeared unquantifiable.

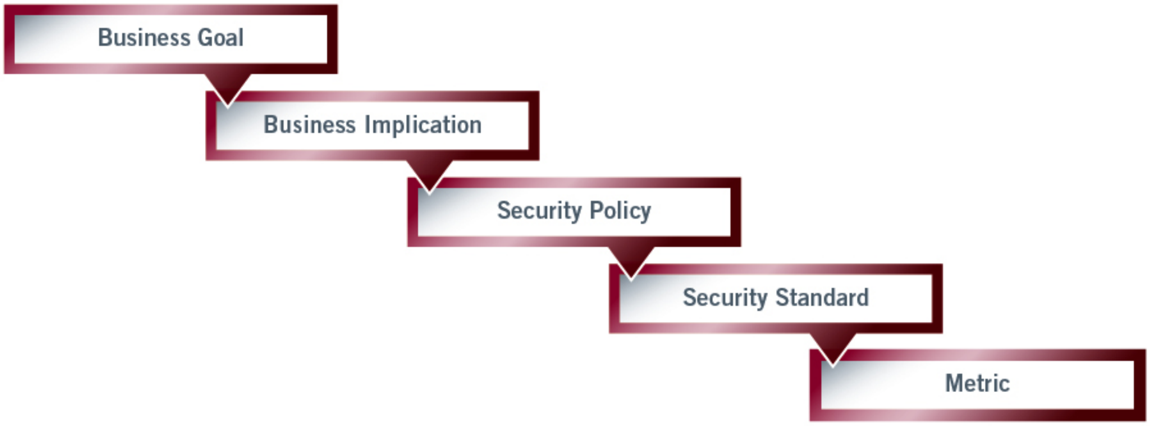

Figure 1: Metrics Derivation

Metrics Progression

If an organization has never collected metrics before, little value will be realized attempting to calculate complex, automated metrics. Figure 2 depicts a natural and strategic metric maturity progression.

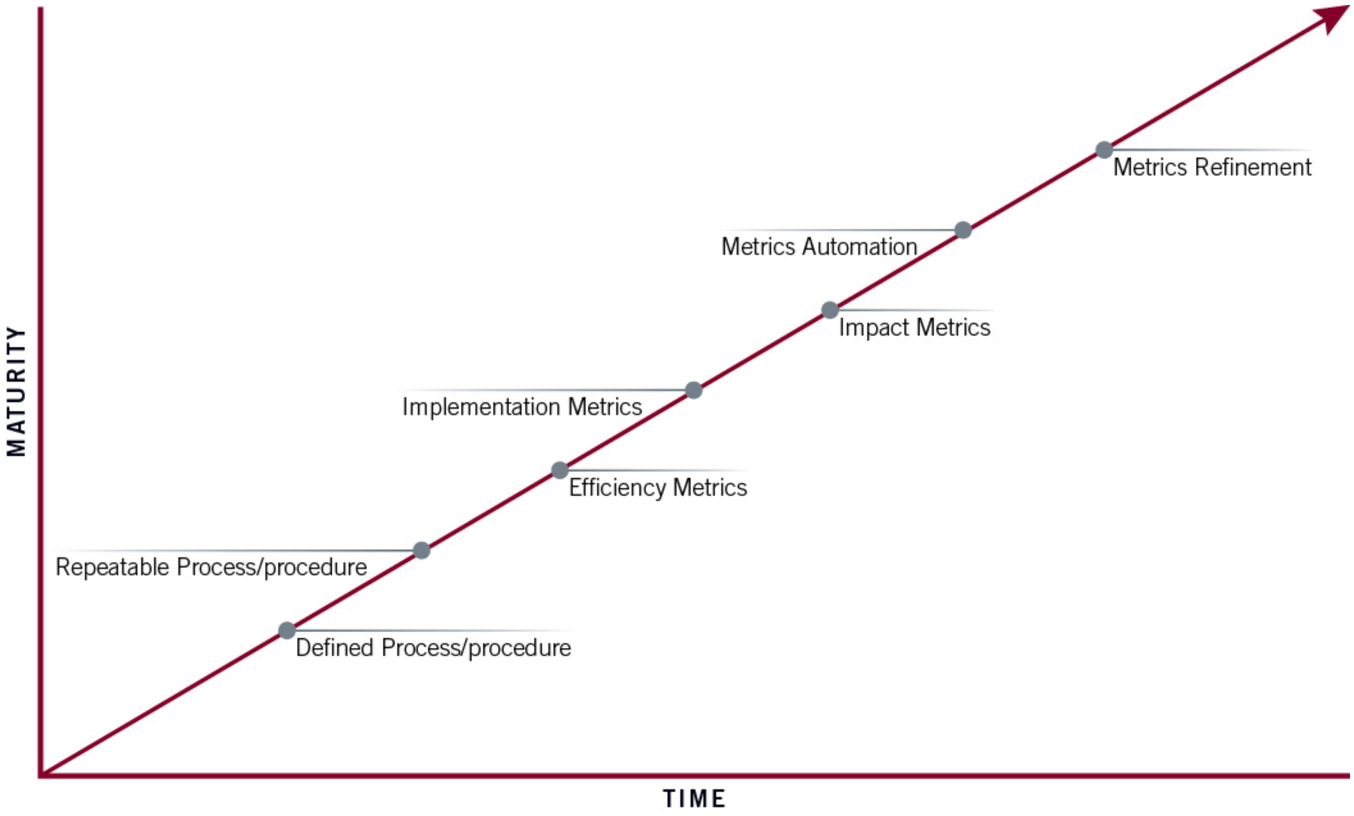

Figure 2: Metric Maturity Progression

A defined and repeatable process or procedure is integral to effective metrics collection. Without this characteristic, the data being captured and calculated is illogically variable. The result of the calculation does not provide a clear analysis, nor does it assist in identifying gaps or areas of strength. On the other hand, a process that is repeatable has such little standard deviation from its normal execution that any variation in the calculated metrics output can be understood as meaningful.

Once mature procedures are established, efficiency metrics can be calculated. DRAIN CVR should be an organization’s primary metric. Once the DRAIN CVR metric has been established, an organization should follow the progression of Figure 2, never attempting to automate a metric before understanding how it works manually. If a CIRT does not believe it can implement the DRAIN CVR metric right away, simplifying the metric and calculating the overall Dwell time and Containment time is recommended, as seen in Figure 3.

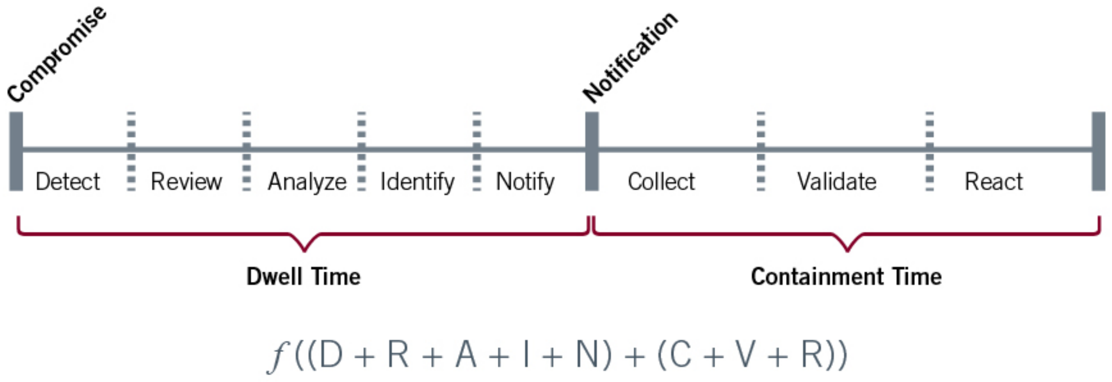

Figure 3: DRAIN CVR Calculation

Finally, a standardized template assists in the definition, collection, dissemination, and review of a metric's output. It is also important to note that when reporting metrics, the measurements should be presented using the most easily digestible means. Where possible, this means reporting results per unit. Very few individuals may understand the importance of a report citing 592,448 critical vulnerabilities over a given reporting time period. However, most audiences will appreciate the same result if delivered as "5.1 critical vulnerabilities per asset."

The table below provides an example that includes several important referential data points regarding a specific metric. These categories are important for efficiency, implementation, and impact metrics.

Field	Data
Metric	01 - Time to Detect
Reference Document	ACME SOP XYZ
Goal	Detect the threat sooner
Measure	Average hours to detect an incident from the moment of initial compromise
Formula	X = (Time to detect incident 1 + time to detect incident 2 + … time to detect incident n)/n
Objective	X < 8 hours
Frequency	Collection: Quarterly Reporting: Quarterly
Trend Retention	Three years
Responsible Stakeholder	CIRT Manager
Data Source	Ticket system / Incident Reports
Reporting Format	Line Graph
Dissemination Mechanism	CIRT Biannual Update

Table 2: Sample Metric Reference Table

CONCLUSION

Metrics are much more than a managerial tool. Utilized to their full potential, metrics are important tools in helping a CIRT more quickly detect, respond to, and contain threats. This increased incident response capability assists an organization in responding to every threat, from commodity malware to The Advanced Persistent Threat. Identifying strengths and weaknesses appropriately must be at the core of every CIRT that wants to maintain or improve its operational capability. Those external to the CIRT often have difficulty understanding the benefits of prevention, or at least of putting a measurable value on those benefits. The proper metrics make this possible and beneficial. Ultimately, a CIRT failing to employ a metrics program will struggle to develop an efficient capability, and will rarely reach a mature capability.

ABOUT MANDIANT SECURITY CONSULTING SERVICES

Security Consulting Services provides long-term assessment, planning, and implementation services for clients to help improve their security posture and effectively prepare for incidents.