

Check Your Blind Spot: A New Cyber-Security Metric for Measuring Incident Response Readiness

Benjamin Aziz¹(✉), Ali Malik¹, and Jeyong Jung²

¹ School of Computing, University of Portsmouth, Portsmouth, UK
{benjamin.aziz, ali.al-bdairi}@port.ac.uk

² Institute of Criminal Justice Studies, University of Portsmouth, Portsmouth, UK
jeyongj@gmail.com

Abstract. This paper presents some ideas on defining and implementing a new Cyber-security risk metric for measuring the readiness of organisations, in terms of the availability of their resources, in dealing with new attack incidents launched against their infrastructures whilst recovering from ongoing incidents. Our new metric, the Mean Blind Spot, is defined as the average interval between the recovery time of an existing incident and the occurrence time of a new incident. It is therefore designed to capture those time intervals where the organisation is most vulnerable due to possible lack of available resources. We present an approach for implementing our new metric using open data on security incidents available from the VERIS community dataset.

1 Introduction

In the context of computing and Cyber systems, measuring risk means choosing an aspect of vulnerability that may exist in a system to investigate, such as its resistance to threats or its exposure to attack incidents. The unit by which risk is measured is usually called the *risk metric*. For example, to measure the frequency at which security attacks occur in some system, one may adopt a risk metric that represents the mean time across these occurrences. Using metrics is a good method for both the quantification of IT risks and reflection of business needs [10]. They are used as objective grounds when an organisation needs to make a decision on its strategy or resource distribution in relation to its IT infrastructure.

Cyber-security risk metrics provide an insight for organisations into the resilience of their IT infrastructure against attacks carried out from over the Internet. As a result, they also give an indication of the cost that may be incurred from the aftermath recovery of such attacks and the cost needed in the future to defend against them. In literature, there have been several efforts that attempt to define and collect such Cyber security-related metrics, examples of which include [4, 7, 11]. And despite recent surveys (e.g. [12]) that question the validity and usefulness of quantified security, we agree more with the view by [3] that

past data are still relevant to new security incidents and that despite the fact that *the road ahead may bend with human whim and technological advance, . . . it does not appear to bend too sharply too often*. Therefore having some idea of the quantitative aspects of security is better than none.

We introduce in this paper a new Cyber-security metric, which we term the *Mean Blind Spot* metric. The new metric is based on the concept of a *blind spot*, which represents the time interval between the moment of occurrence of a new security incident and the moment at which an existing incident has been fully recovered. As such, a blind spot reflects the notion of *readiness* of an organisation or its IT security team to deal with new security incidents as they occur while dealing with the recovery from existing ones. Such readiness assumes that the deployment of resources to the recovery of incidents can only contribute positively to that recovery. Although our new metric does not identify the cause of a problem nor suggest a solution for the cause, it can work as objective evidence when an IT manager argues for more organisational support or resources to secure their infrastructure. We define an implementation of this metric in an open-source community dataset.

The rest of the paper is structured as follows. In the next Sect. 2, we review related work including the collection of Cyber-security and network security risk metrics defined in [4, 7, 11]. In Sect. 3, we give a quick background on a couple of closely related Cyber-security metrics and demonstrate their definitions with a simple running example. In Sect. 4, we introduce our new metric, the Blind Spot, and discuss its rationale and definition, including some variations that represent higher-level views of the problem of blind spots. In Sect. 5, we present our implementation using the VERIS dataset. Finally, in Sect. 6, we conclude the paper and give directions for future work.

2 Related Work

As the dependence on ICTs of an organisation increases, it is important that information security is integrated into business strategies. Many studies [5, 6, 9] suggest that senior management should discuss IT agendas and issues as business matters. However, top decision-makers are not familiar with IT terms but with business language. If information security issues are not explained in business terms, it may be hard to gain support from senior managers. Generally, there needs to be two things in place to aid the understanding of senior managers. The first is the quantification of IT issues and responsive measures. When security risks and countermeasures are quantified, it becomes much easier to calculate business impact resulting from IT issues. The second relates IT issues to business goals and objectives, where IT agendas reflect needs of businesses [5].

Many attempts have been made to suggest standardised Cyber-security metrics for organisations. Each study has a different approach. As an international body of the UN, the Telecommunication Standardisation sector of the International Telecommunication Union (ITU-T) published Cyber-security indicators of risk [11], which included not only technical factors but also human factors

as well. Indicators such as “security training and education” and “personnel security” were adopted to reduce human errors or intended behaviours in an organisation.

On the other hand, the Center for Internet Security (CIS) metrics [4] focus mostly on technical and business factors without consideration of human factors. The CIS defined in [4] seven metrics that are directly related to the overall incident management process, ranging from incident detection to incident recovery. We adopt two such technical metrics defined in [4] as the basis for our work here.

Criticising past metrics as “labour intensive” and “subjective”, Lippman et al. [7] argued that continuous risk assessment based on a data-driven approach was necessary to reflect the constantly changing nature of threats. The metrics proposed in [7] are of complex mathematical nature and hence their applicability is questionable. Chew et al. [2] suggest three types of metrics used differently depending on the purpose and nature of a metric. Implementation metrics are intended to measure the extent at which security policies are implemented. Secondly, effectiveness/efficiency metrics measure how well security services are delivered. Lastly, impact metrics aim to measure impacts of security incidents on a business.

One could argue that the work presented here involves the second type of metrics, since the aim of the work is to define metrics that measure the readiness of an IT department within an organisation when facing incidents over time. Measuring the readiness of security services allows for the diagnoses of an organisation on its capability of handling unexpected incidents.

Payne [8] suggested seven key steps to establishing a security metrics programme. One of them is to establish benchmarks and targets. Setting benchmarks is useful when evaluating success or failure of current security controls [1]. There should be some criteria for benchmarks. Too simplistic metrics may not be appropriate for being regarded as benchmarks because they are naturally intuitive or self-explanatory. Thus, creating an advanced metric based on basic ones is a good practice that we adopt in our approach. Also, metrics for benchmarks need to be used for driving improvements for existing practices. It means that they have actual impact on IT or business management.

In our case, we adopt a widely used large community dataset called VERIS [13] as our benchmark on which we implement our new incident readiness metrics. After a benchmark is adopted in an organisation, there is no hard and fast rule as to choosing a reference point for the benchmark. The choice of acceptable levels for our metrics will depend on organisational context.

3 Background

Literature has numerous metrics related to Cyber security (e.g. [4, 7, 11]). We give here an overview of two such closely related metrics defined in [4], which we use later as part of the definition of our new set of metrics. We also give an overview of a widely-used security incident vocabulary and dataset known as VERIS, which we use as a benchmark reference for the implementation of our new metrics.

3.1 Mean Time Between Security Incidents

The Mean Time between Security Incidents (MTBSI) metric is described in [4] as a metric for calculating the mean time between occurrences of security incidents in some organisation's IT infrastructure. This type of operational metrics can be defined by the following formula:

$$MTBSI = \left(\sum_{i=1}^{n-1} (Date_of_Occurrence(incident_{i+1}) - Date_of_Occurrence(incident_i)) \right) / (n - 1) \quad (1)$$

where n is the total number of recorded incidents. As a result, there would be only $n - 1$ intervals between any n incidents. We consider the unit of measurement of the MTBSI metric to be time, e.g. hours, days, weeks etc. The following Table 1 shows an example of 10 incidents recorded with the dates and times of their occurrences.

Table 1. An example of incident occurrence dates and times

Incident number	1	2	3	4	5
Date of occurrence	01.06	01.06	01.06	01.06	01.06
Time of occurrence	12:10	12:50	14:00	14:56	18:30
Incident number	6	7	8	9	10
Date of occurrence	01.06	02.06	02.06	02.06	02.06
Time of occurrence	18:35	07:20	09:20	12:30	19:40

To calculate the MTBSI for this example, we evaluate Eq. (1) above:

$$MTBSI = \frac{(40+70+56+214+5+765+120+190+430)}{9} = 210 \text{ min.}$$

This means that, on average, there are 3.5 h separating the occurrence of any two incidents.

3.2 Mean Time to Incident Recovery

The second widely-used metric for measuring Cyber security is the Mean Time to Incident Recovery (MTIR), which reflects the mean time needed from the moment an incident occurs to the moment it is recovered.

This type of operational metrics can be defined using the following formula from [4]:

$$MTIR = \left(\sum_{i=1}^n (Date_of_Recovery(incident_i) - Date_of_Occurrence(incident_i)) \right) / n \quad (2)$$

where n is the total number of recorded incidents. We take the unit of measurement for MTIR again to be time, e.g. hours, days, weeks etc. Note that we divide over n since the number of recoveries is the same as the number of incidents occurring. For example, in the following Table 2, we have again the same 10 incidents recorded from Table 1, but this time also with their dates and times of recovery.

Table 2. An example of incident occurrence/recovery dates and times

Incident number	1	2	3	4	5	6	7	8	9	10
Date of occurrence	01.06	01.06	01.06	01.06	01.06	01.06	02.06	02.06	02.06	02.06
Time of occurrence	12:10	12:50	14:00	14:56	18:30	18:35	07:20	09:20	12:30	19:40
Date of recovery	01.06	01.06	01.06	01.06	01.06	01.06	02.06	02.06	02.06	03.06
Time of recovery	13:55	14:40	19:30	19:05	20:10	21:30	11:10	13:50	15:50	00:15

To calculate MTIR for this example, we evaluate Eq. (2) above:

$$MTIR = \frac{(105+110+330+249+100+175+230+270+200+275)}{10} = 204.4 \text{ min.}$$

This means that each incident takes on average about 3 h and 24 min to recover.

3.3 VERIS

The Vocabulary for Event Recording and Incident Sharing (VERIS) [13] is a dataset and schema capturing a set of metrics for describing security incidents. It is currently considered a leading provider of open quality information in the IT security domain and provides a framework that organisations can use to collect and share information on security incidents in a responsible and anonymous manner, with the aim of constructing a ground on which researchers and experts in the IT security industry can cooperate to learn from their experiences. We use the dataset provided in VERIS, known as VCDB [14], as a benchmark on which we implement our new blind spot-based metrics defined in the next sections.

The VERIS schema itself consists of five general sections, containing descriptions of the security incidents in the VERIS dataset. These five categories are as follows:

- *Incident Tracking*: this section contains general information about the incidents, for example, the source identity, summary of the incident and whether the incident is related to other incidents.
- *Victim demographics*: this section contains information related to the organisation being affected by the incident, for example, its country of operation, number of employees, revenue and industry type.
- *Incident description*: this section contains information related to the question of “who did what to what (or whom) with what result”.

- *Discovery and response*: this section contains information related to the incident’s timeline, its discovery method, root causes, corrective actions etc.
- *Impact assessment*: this last section contains information on loss categorisation and estimation, impact rating and so on.

For the purpose of this paper, we are mainly interested in one kind of information; namely *time to containment*. This is the closest in nature to the MTIR metric described above, and appears under the “Discovery and response” section of information. In VCDB, this metadata appears as *timeline.containment*. The available meaningful values for the timeline unit for this metadata include seconds, minutes, hours, days, weeks, months, years and never. Other values are NA and unknown, but we do not consider these to be useful.

The significance of the VERIS dataset lies in the fact that it is a *community-based* dataset. This means that its data are collected from a wide range of industries and varied over different types and sizes of organisations. This renders it more interesting and with wider applicability than datasets generated in single organisations.

4 The Mean Blind Spot Metric

Our new incident readiness metrics rely on a concept we call the *Blind Spot* (BS). A BS is the time interval between the moment a new security incident occurs and the last moment the previous security incident was recovered, as shown in Fig. 1. In its worse case, a BS represents the time when an organisation has to start recovery from a new incident whilst still recovering from an earlier one. We consider this metric to be an indication to the readiness of an organisation to encounter new incidents and a measure of the vulnerability organisations may face in such situations where not enough resources are available to recover from security incidents.

Note at this stage that, for the sake of simplicity, we do not consider part of this model the scenario when two incidents arrive exactly at the same moment in

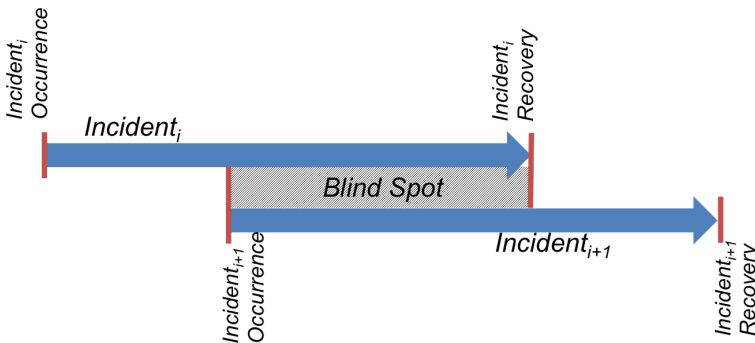


Fig. 1. A Blind Spot

time (i.e. when $Date_of_Occurence(incident_i) = Date_of_Occurence(incident_{i+1})$). This is justified since later during the VERIS-based implementation part, we replace this difference in arrival time with the MTBSI metric (and again assume that $MTBSI > 0$).

We can average out this difference in occurrence times of new incidents and the recovery times of older ones in terms of the *Mean Blind Spot* (MBS) metric as follows:

$$MBS = \left(\sum_{i=1}^{n-1} (Date_of_Occurence(incident_{i+1}) - Date_of_Recovery(incident_i)) \right) / (n - 1) \quad (3)$$

The unit of measurement for the MBS metric is time, e.g. hours, days, weeks etc. The mean is calculated over $n - 1$, as there are only $n - 1$ blind spots for n number of recorded incidents, as shown in Fig. 2.




Fig. 2. $n - 1$ Chain of Blind Spots

There are a couple of important assumptions this definition relies on:

- The definition assumes a *first-come-first-serve* model of scheduling incidents to recovery resources, or in other words, incident i is scheduled for recovery before incident $i + 1$. This is important for the blind spot time area to be a true one, otherwise it will contain the idle time that an incident spends waiting in the scheduling queue.
- The second assumption relies on the fact that all the recovery resources will contribute positively to all the occurring security incidents. In reality, this may not always be the case. Some resources may require some time to become positive contributors to the reduction of the recovery time for an incident. In fact, some resources may only have negative contribution to the recovery of an incident. This situation is depicted in Fig. 3, where we only assume Type 1 resources in our model.

Let’s consider how the MBS metric works through an example. The following Table 3 shows again our 10 security incidents with their occurrence and recovery times, but this including also their blind spot times.

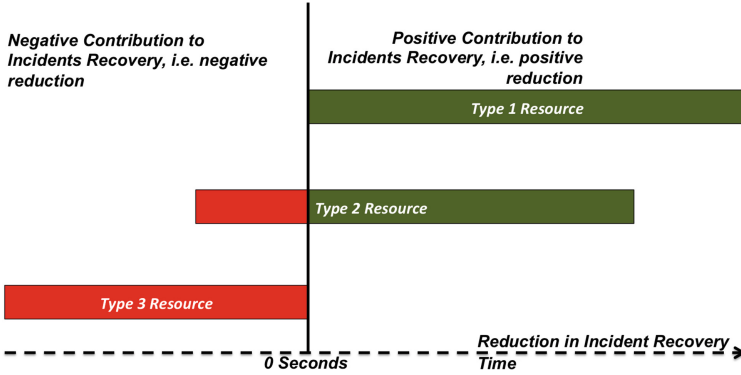


Fig. 3. Resources Contribution to Incident Recovery

Table 3. An example showing blind spots

Incident number	1	2	3	4	5	6	7	8	9	10
Occurrence date	01.06	01.06	01.06	01.06	01.06	01.06	02.06	02.06	02.06	02.06
Occurrence time	12:10	12:50	14:00	14:56	18:30	18:35	07:20	09:20	12:30	19:40
Recovery date	01.06	01.06	01.06	01.06	01.06	01.06	02.06	02.06	02.06	03.06
Recovery time	13:55	14:40	19:30	19:05	20:10	21:30	11:10	13:50	15:50	00:15
Blind spot interval	-65	-40	-274	-35	-95	590	-110	-80	230	-

Note that for the last incident, there is no blind spot time as no incidents are recorded after that. For this example, we can calculate MBS as follows:

$$MBS = (((-65) + (-40) + (-274) + (-35) + (-95) + 590 + (-110) + (-80) + 230))/9 = 13.44 \text{ min.}$$

A positive value (as in this case) for the MBS metric is good, since it indicates that there is a positive time margin between, on average, the occurrence and recovery of incidents. However, a negative value would signal no such margin exists and that incidents’ recovery stages are overlapping. This may further have implications on an organisation’s capability to cope with the speed of occurrence of security incidents since recovery from earlier incidents is, on average, slow.

We next discuss one variant of this metric, which incorporates an organisation’s appetite for blind spots.

4.1 An Approximated MBS

A first variation of the MBS metric that we introduce is an approximated one, which can be calculated directly using the MTIR and MTBSI metrics discussed earlier. We call this variation the Approximate MBS (AMBS) metric. The general formula for the AMBS metric is as follows:

$$AMBS = \frac{MTIR}{MTBSI} \quad (4)$$

The intuition behind this metric is that it gives some sense of how large the difference is between recovery and incident occurrence intervals as captured by the MTIR and MTBSI metrics, respectively. Therefore, it provides a quick way of understanding the effect of the blind spot problem. If, on average, recovery intervals are smaller than incident occurrence intervals, then this ratio would be less than one, which is good for the organisation. If, on the other hand, the ratio is one or more, it means that the occurrence intervals are at least as large as the recovery ones, on average, which is bad for the organisation.

Consider again the example of the previous section. We calculated that $MTIR = 204.4$ min/interval and that $MTBSI = 210$. Therefore, one can calculate $AMBS = 204.4/210 = 0.973$. This value, enforces the conclusion arrived at by the calculation of the MBS metric that on average, in the case of our example, blind spots do not pose a problem in terms of overall time they last. As we see later in Sect. 5, this metric also gives an indication as to the maximum number of incidents an organisation may be recovering from in any one moment in time.

4.2 Ratio of Blind Spots Metric

The Ratio of Blind Spots (RBS) metric is not, strictly speaking, based on the MBS metric but more fundamentally based on the concept of a blind spot. In order to define RBS, we first define a Blind Spot Appetite (BSA) value, which represents the maximum blind spot time an organisation or an IT team is willing to tolerate. For example, a BSA value might be -60 min, meaning that the organisation is willing to tolerate scenarios where recovery from an existing incident overlaps the occurrence of a new one in a maximum of one hour.

The ratio BS/BSA therefore represents a measure of how far a blind spot is from the appetite value. A value of $BS/BSA = 1$ or less means that the blind spot is within the acceptable range and a value of more than 1 means that the blind spot is unacceptable. For simplicity, we approximate all the values of $BS/BSA < 0$ to 0, since in this case these have the same meaning as to when $BS/BSA = 0$. Returning to the example of the previous section, we calculate the BS/BSA values for each blind spot as shown in Table 4.

Based on the BS/BSA ratio, one can define the new RBS metric as follows, in terms of the cardinality of a multiset (bag) of all those ratios who's value is over 1:

Table 4. Example showing the ratio BS/BSA

Incident number	1	2	3	4	5	6	7	8	9	10
Blind Spot interval	-65	-40	-274	-35	-95	590	-110	-80	230	
BS/BSA ratio	1.08	0.67	4.57	0.58	1.58	0	1.83	1.33	0	

$$RBS = \frac{\text{card}(\{y \text{ where } (y = BS/BSA) \wedge (y > 1)\})}{n - 1} \times 100\% \quad (5)$$

The RBS metric hence captures the percentage of the ratio of all the BS/BSA elements, which are over 1, over the overall number of blind spots. Unlike MBS, it does not rely on a mean-based calculation, but represents more the percentage of “risky” blind spots in an organisation or an IT team. In our example, $RBS = \text{card}(\{1.08, 4.57, 1.58, 1.83, 1.33\})/9 \times 100\% = 56\%$. This means that, despite the fact that MBS is on average positive, 56% of blind spots are risky.

5 Method Implementation Using VERIS

In this section, we propose a practical approach for implementing our new metrics using the VERIS dataset (VCDB) [14]. This implementation will allow organisations to obtain some idea of their level of readiness in dealing with blind spots, without the need for much precise information about their own security incidents.

5.1 Implementing the MBS Metric

Our first implementation provides a measurement function for new organisations to assess their level of readiness based on two pieces of information: First their MTBSI metric values and second the time to containment metric in the VERIS dataset. Note that here we parameterise by MTBSI since VERIS, despite its rich collection of incident metadata, does not specify whether two incidents belong to the same organisation and in what temporal order they occur.

We define the signature of the blind spot readiness measurement function, f , as follows:

$$f : \text{Time} \rightarrow \text{Percentage} \quad (6)$$

which takes in a time unit expressing the MTBSI for the particular organisation, and returns a percentage number expressing the level of blind spot readiness for that organisation. This is the compliment of the percentage of incidents that are deemed to be *risky* with respect to the information provided by the VERIS dataset, in the sense that there is high likelihood that the organisation may not be prepared to contain them in good time.

Our definition of f is constrained by two aspects of the VERIS dataset: First, there are no timeline information across the reported incidents, which means that it is not possible to conclude, given two incidents, what their sequence is. Second, no concrete timeline data is given; only time units (e.g. hours, days, weeks, etc.) As a result, our implementation relies on the relationship between the lengths of the MTIR (i.e. time to containment) and MTBSI metrics when deciding whether a blind spot exists or not.

Figure 4 depicts the relationship between a blind spot and the MTIR and MTBSI metrics. In the absence of concrete dates/times marking the start and

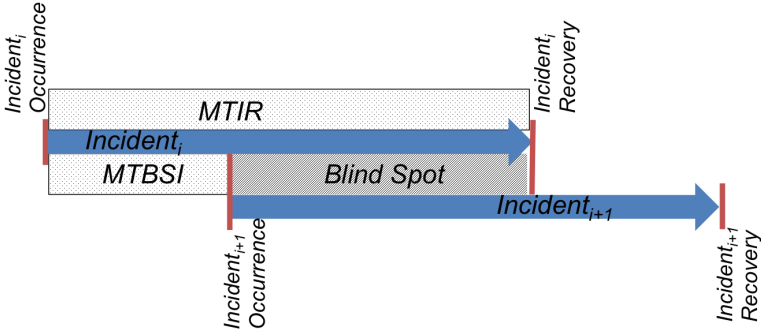


Fig. 4. Relationship between a Blind Spot and MTIR, MTBSI

recovery points of incidents, we consider MTBSI as the metric describing the uniform time difference between consecutive incident occurrences, and MTIR as the metric describing the uniform time between the start and recovery times of incidents. As a result, an MTIR value that extends beyond MTBSI is in a blind spot area, and one that does not is not. As a consequence of the lack of precise information on incident occurrence and recovery times in VERIS, one can only implement MBS as the difference between MTBSI and MTIR (i.e. $implementation(MBS) = MTBSI - MTIR$). This is reasonable since the definition of a blind spot between two incidents i and $i + 1$ is such that $Date_of_Occurrence(incident_{i+1}) - Date_of_Recovery(incident_i)$. However, we can arrive at this by performing $(Date_of_Occurrence(incident_{i+1}) - Date_of_Occurrence(incident_i)) - (Date_of_Recovery(incident_i) - Date_of_Occurrence(incident_i))$, which is the difference between MTBSI and MTIR, assuming a uniform value for all incidents. Therefore, our implementation function f relies on this difference, and can be defined in the following manner:

$$f(MTBSI) = (100\% - \sum_{i=1}^n (percentage(c_i))) \text{ where } c_i \geq time_unit(MTBSI)$$

where c represents the time unit (i.e. seconds, minutes, hours, days etc.) for the Discovery-to-Containment stage in the timeline of events, and hence $percentage(c)$ is the percentage of all incidents where the time to containment metric has been reported to be in that specific time unit. On the other hand, n represents the number of time units that are larger or equal to the MTBSI's time unit, as returned by the auxiliary function $time_unit$. For example, if $MTBSI = 15$ hours/incident interval, then $time_unit(MTBSI) = hours$ and $n = 6$, where the six time units in this case would be $\{hours, days, weeks, months, years, never\}$. We enumerate these as c_1, \dots, c_6 . Note here that we also include the hours time unit, in order to err on the safe side. We also exclude those incidents with a "NA" or "unknown" values. Finally, i ranges over n .

Table 5. Percentages of incidents for each time unit of the time to containment metric (VERIS, 2013 Quarter 4)

Time unit	Seconds	Minutes	Hours	Days	Weeks	Months	Years	Never
Percentage (%)	2.17	5.07	42.03	29.00	7.97	7.97	2.17	3.62

Considering the 2013Q4 version of the dataset, we have the following percentages of incidents for each time unit of the time to containment metric, as shown in Table 5.

These numbers are based on a dataset size of 2476 incidents. Going back to the example above of MTBSI = 15 hours/incident interval, one can calculate f as follows:

$$\begin{aligned}
 f(15 \text{ hours/incident interval}) &= (100 - \sum_{i=1}^6 (\text{percentage}(c_i)))\% \\
 &\text{where } c_i \geq \text{time_unit}(15 \text{ hours/incident interval}) \\
 &= (100 - (\text{percentage}(\text{hours}) + \text{percentage}(\text{days}) + \\
 &\quad \text{percentage}(\text{weeks}) + \text{percentage}(\text{months}) + \\
 &\quad \text{percentage}(\text{years}) + \text{percentage}(\text{never})))\% \\
 &= (100 - (42.03 + 29 + 7.97 + 7.97 + 2.17 + 3.62))\% \\
 &= 7.24\%
 \end{aligned}$$

This means that the organisation, according to its reported MTBSI value of 15 hours/incident interval, will only be fully ready in 7.24% cases of security incidents based on the data provided in the VERIS dataset. In 92.76% of cases, the organisation may/would struggle to cope with new security incidents according to the blind spot readiness metric. On the other hand, if for example the MTBSI was 15 weeks instead, then the above value returned by f would rise to 78.27%.

5.2 Implementing the AMBS Metric

The second implementation we introduce will simply be an implementation of the AMBS metric. Recall that the AMBS metric is simply dividing the length of the MTIR metric by the length of the MTBSI metric, as depicted in Fig. 5.

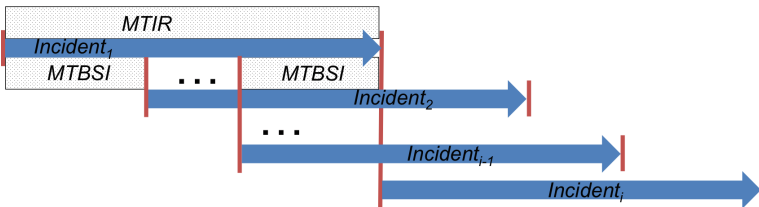


Fig. 5. A depiction of the ratio of MTIR to MTBSI

In addition to showing the ratio of the two metrics, it turns out that this definition can be used to estimate the minimum and maximum number of incidents that an organisation will have to deal with at any one time. One can deduce this fact from considering that the start of every MTBSI period signals the start of a new security incident. Therefore, if one was to fix a time frame within which one could count the number of MTBSI periods, then this would also imply the number of incidents within that time frame.

The latter case of the maximum number of incidents is particularly of interest from an incident readiness point of view, as it provides the management team some idea of the scale of resources required to tackle such events. The minimum number, on the other hand, will provide an indication of what level of resource relaxation the organisation can reach.

We start first by defining the following auxiliary function, f_{aux2} :

$$f_{aux2}(MTBSI, c) = \{\lceil \min(c)/MTBSI \rceil, \lceil \max(c)/MTBSI \rceil\}$$

which returns a multiset (bag) of two elements. These elements are the minimum and maximum number of incidents the organisation will be recovering from at any one time, corresponding to its specific value of MTBSI and the time unit, c , of the *time-to-containment* metric (i.e. the MTIR metric) as defined in VERIS. The MTBSI value is necessary here, so this has to be supplied by the organisation. However, the time unit c is ranged over all the meaningful time units defined in VERIS (i.e. days, hours, weeks etc.).

We calculate these numbers as the ceiling (“the gallows”) ratio between the minimum and maximum values we approximate for c as explained below and the supplied value for MTBSI. This is needed since we consider that a fraction of an incident is safer approximated to a whole incident (i.e. next integer up).

We next explain how $\min(c)$ and $\max(c)$ are defined. Since c itself is only a time unit due to the lack of concrete date/time information on VERIS-recorded incidents, we require a 2-point time value concretisation of this abstract time unit. We do this based on the following ranges (assuming a month is 4.35 weeks):

$$\begin{aligned} \text{range(seconds)} &= [1 \text{ s}, 60 \text{ s}) \\ \text{range(minutes)} &= [1 \text{ min}, 60 \text{ min}) \\ \text{range(hours)} &= [1 \text{ h}, 24 \text{ h}) \\ \text{range(days)} &= [1 \text{ day}, 7 \text{ days}) \\ \text{range(weeks)} &= [1 \text{ week}, 4.35 \text{ weeks}) \\ \text{range(months)} &= [1 \text{ month}, 12 \text{ months}) \\ \text{range(years)} &= [1 \text{ year}, \infty \text{ years}) \\ \text{range(never)} &= [\infty \text{ years}, \infty \text{ years}] \end{aligned}$$

Note that we do not consider the case of “never”, since this will result in dividing ∞ by MTBSI, which returns an infinite number of incidents. Since c effectively provides a *time window* within which the maximum and minimum number of incidents are approximated, providing an infinite time window will naturally lead to an infinite number of minimum and maximum incidents. We do not consider this information meaningful for understanding an organisation’s

readiness in tackling those periods of time when number of incidents is at its maximum.

Let's consider now an example of how f_{aux2} works. We assume that we are calculating the function for the case of $c = \text{Weeks}$, then we need to consider the following two points in time: At week 1 and at week 4.35. This is because any less than 1 week the metric would turn to a daily time unit and any more than 4.35 weeks the metric would turn into months. If the organisation provides a value for $MTBSI = 18 \text{ days} = 2.57$, then this means that

$$f_{aux2}(18 \text{ Days}, \text{Weeks}) = \{\lceil 1/2.57 \rceil, \lceil 4.35/2.57 \rceil\} = \{1, 2\}$$

On the other hand, if $MTBSI$ was to drop to 1 day (i.e. 0.143 week), then the number of incidents would increase:

$$f_{aux2}(1 \text{ Day}, \text{Weeks}) = \{\lceil 1/0.143 \rceil, \lceil 4.35/0.143 \rceil\} = \{7, 31\}$$

Our function then for generating an AMBS estimation from the VERIS dataset would pair each of these two values with the percentage of incidents the c time unit occurs in VCDB:

$$f_2(MTBSI, c) = (f_{aux2}(MTBSI, c), \text{percentage}(c)) \quad (7)$$

Therefore, for the case of $c = \text{Weeks}$, we have that $\text{percentage}(\text{Weeks}) = 7.97\%$ from [13], and hence $f_2(18 \text{ Days}, \text{Weeks}) = (\{1, 2\}, 7.97\%)$ whilst for the case of $f_2(1 \text{ Day}, \text{Weeks}) = (\{7, 31\}, 7.97\%)$. The meaning of these pairs is to provide an approximate percentage (in this case 7.97%) of the likelihood of the maximum/minimum incident number estimations being true, with reference to the data provided in VCDB and the selected c time unit. If the selected c time unit was changed, say to Days, the pair would become (for the case of $MTBSI = 18 \text{ days}$) $(\{1, 1\}, 29\%)$.

6 Conclusion and Future Work

We presented in this paper a new risk metric, called the blind spot, for expressing Cyber incident recovery readiness in organisations and IT departments. The new metric represents the gap in time between the recovery (or containment) of existing incidents and the occurrence of new incidents. We postulate that the longer the gap, the more vulnerable the organisation or IT department will be to lack of resources in tackling new incidents, hence the relationship with the concept of incident response readiness. Furthermore, we defined three variants of this new metric: the Mean Blind Spot, the Approximated Mean Blind Spot and the Ratio of Blind Spots metrics. We demonstrated how these metrics can be implemented over an open source large dataset containing information on Cyber security incidents, namely the VERIS dataset.

The significance of the VERIS dataset lies in the fact that it is community-driven, where data are collected from a variety of organisations in a wide range of industries covering small, medium and large size organisations. This ensures

that the implementation of the new metrics is applicable to the wider community. However, we plan in the future to further validate the new metrics based on empirical data obtained from specific case studies for IT teams and organisations. Such specific case studies produce more accurate results, despite their scope of applicability. The application of real empirical data to this metric may expose more (specific) benefits and drawbacks for the new metrics, possibly suggesting ways to refine our initial conceptual model. Such studies will also help incorporate new factors or new metadata into the current model, particularly since the characteristics of Cyber security incidents vary over time also depending on the context. Therefore, the refinement of the metric can yield more benefits for organisations.

References

1. Black, P.E., Scarfone, K., Souppaya, M.: Cyber security metrics and measures. In: Voeller, J.G. (ed.) *Wiley Handbook of Science and Technology for Homeland Security*, Chap. 5, pp. 1–15. Wiley, London (2008)
2. Chew, E., Swanson, M., Stine, K., Bartol, N., Brown, A., Robinson, W.: *Performance measurement guide for information security*. Technical report 800–55 Revision 1, National Institute of Standards and Technology, July 2008
3. Hoo, K.J.S.: *How Much is Enough? A Risk-Management Approach to Computer Security* (2000)
4. The Center for Internet Security: *CIS Security Metrics v1.1.0*, November 2010
5. Kayworth, T., Whitten, D.: Effective information security requires a balance of social and technology factors. *MIS Q. Executive* **9**(3) (2012). <http://ssrn.com/abstract=2058035>
6. Kwon, J., Ulmer, J.R., Wang, T.: The association between top management involvement and compensation and information security breaches. *J. Inf. Syst.* **27**(1), 219–236 (2013). <http://dx.doi.org/10.2308/isis-50339>
7. P-Lippmann, R., Riordan, J.F., Yu, T.H., Watson, K.K.: *Continuous security metrics for prevalent network threats: introduction and first four metrics*. Technical report ESC-TR-2010-099, Massachusetts Institute of Technology (2012)
8. Payne, S.C.: *A guide to security metrics*. Technical report SANS Security Essentials GSEC Practical Assignment, Version 1.2e, Escal Institute of Advanced Technologies, Inc. (The SANS Institute), June 2006
9. von Solms, B., von Solms, R.: From information security to. business security? *Comput. Secur.* **24**(4), 271–273 (2005)
10. Swanson, M., Bartol, N., Sabato, J., Hash, J., Graffo, L.: *Security metrics guide for information technology systems*. Technical report 800–55, National Institute of Standards and Technology, July 2003
11. International Telecommunication Union: *A Cybersecurity indicator of risk to enhance confidence and security in the use of telecommunication/information and communication technologies*. Technical report X.1208, International Telecommunication Union (2014)
12. Verendel, V.: Quantified security is a weak hypothesis: a critical survey of results and assumptions. In: *Proceedings of the 2009 Workshop on New Security Paradigms Workshop, NSPW 2009*, pp. 37–50. ACM, New York (2009)
13. VERIZON: *The Vocabulary for Event Recording and Incident Sharing (VERIS)*. <http://veriscommunity.net/>, Accessed 21 Nov 2016
14. VERIZON: *VERIS Community Database*. <http://vcdb.org/>, Accessed 21 Nov 2016