

Reporting Cyber Risk to Boards

CISO Edition

Control, Measure, Report, Repeat

Authors

Freddy Dezeure

George Webster

Jason Trost

Eireann Leverett

João Pedro Gonçalves

Patrick Mana

Greg McCord

Josh Magri

Reviewers

Lokke Moerel

Alex Iftimie

Chris Deverell

Greg Bell

Jamie Hutchinson

Date: 28 March 2022

Version: Final

Table of Contents

INTRODUCTION	3
PURPOSE	4
BOARD EXPECTATIONS	4
IT'S ALL ABOUT RISK	5
IMPORTANT CHOICES TO MAKE	7
QUANTITATIVE METRICS.....	9
A METRICS MODEL.....	11
GATHERING INPUTS – MEASURE WHAT MATTERS MOST	12
DATA SOURCES – UNDENIABLE TRUTH.....	18
TRANSFORMING – ARE OUR CONTROLS GOOD ENOUGH?	18
REPORTING CYBER RISK –PROVIDE REASONABLE ASSURANCE	23
COMMUNICATION CHANNEL(S)	24
OVERCOMING RESISTANCE	25
ALLOCATING RESOURCES TO METRICS AND REPORTING	26
ANNEX 1: WHERE TO START?	27
ANNEX 2: EXAMPLES FROM THE COMMUNITY.....	28
ANNEX 3: SAMPLE REPORT	33

Introduction

This document provides methods and inspiration for Chief Information Security Officers (CISO) to design and implement quantitative cybersecurity metrics to report cyber risk at Board level and provide reasonable assurance that the risk is within the accepted risk appetite.

Once upon a time, you could protect your secrets by turning a key in a closed door. For your deepest secrets, you might have installed a better door, maybe improved the walls, or stationed a couple of guards. When you needed to move your secrets, you would bundle them into a bag and use steganography or cryptography to keep the secret from prying eyes. This fairy tale was true for computers too, but this time is long gone. Our society, economy, and day-to-day life depend on the exchange of information that swims through our interconnected systems. The concept of a protective fence is a thing of the past.

The modern economy and its reliance on data has made our secrets ever growing in value, and this has attracted the attention of the professional criminal, ever probing our defences. Our information systems create substantial risk to governments, businesses, and individuals alike. In 2021, \$4 million was the average cost for a data breach at a typical corporation. A major breach could even go upwards of \$400 million¹. The total costs for all cybersecurity incidents in 2020 are estimated at \$1t, a more than 50% increase in two years².

It is no wonder that cybersecurity is a top-of-mind issue for most organizations and governments, and this attention is rightfully deserved. As an example, the new SEC regulations related to cybersecurity risk disclosures include provisions on the importance of communicating cybersecurity risk to boards³.

But having the ear of senior stakeholders is not solving the cybersecurity problems or reducing the risk. Our business and governmental leaders are ill-equipped to deal with cybersecurity because cybersecurity does not speak their language. In turn, cybersecurity is ill-equipped to deal with senior stakeholders because cyber professionals struggle to measure their program's effectiveness, articulate program utility, or even communicate its successes. This inability to measure the effectiveness of cybersecurity controls and communicate the risk reduction they produce to the senior stakeholders puts cybersecurity professionals in a position where they jockey for budget, yet they do not know if what they are doing is actually reducing the risk of losses.

Zero risk is unreachable and unrealistic. It always has been. But the dynamics have changed. Furthermore, the pace of change in the cybersecurity threat landscape exceeds our ability to adjust controls for risk or even identify which controls matter for mitigating risk. A CISO needs to justify the cybersecurity budget and explain why the chosen approaches align with the organization's overall risk appetite. This is a challenging task, but one for which this paper aims to provide inspiration and material to design and implement pragmatic solutions.

¹ <https://www.ibm.com/security/data-breach>

² <https://www.mcafee.com/enterprise/en-us/assets/reports/rp-hidden-costs-of-cybercrime.pdf>

³ <https://www.mofo.com/resources/insights/220311-sec-proposes-cybersecurity-disclosure-rules.html>

Purpose

This paper presents orientations for CISOs to report cyber risk and its context to their senior stakeholders, such as their Board. It describes methods that help CISOs engage in cyber risk management, communicate this effectively, and facilitate proper oversight. While not a focus of this paper, the content of this document also helps with reporting cyber risk to other stakeholders, like regulators, insurers, and clients.

We believe that metrics are a necessary component of any successful effort. Peter Drucker's management theories are a testament to this and changed business as a result. Cybersecurity is no different. Metrics are needed to guide the management of risks in organizations, to increase cyber resilience over time, and to show compliance to internal and external stakeholders. However, you must measure the right things and look at the problem holistically. Unfortunately, measuring cybersecurity risks and applying an appropriate metric is akin to finding the holy grail. Good practices for metrics are few, their dissemination across the community is rare, and measuring risk in a non-deterministic cybersecurity threat landscape is challenging.

This paper summarizes findings and best practices from a CISO Working Group. Full credit is due to the participants. Without their insights, sharing and interaction, this paper would not have been possible. Included within this paper are annexes with examples from the community. We hope to generate additional contributions so that a new body of work and example implementations emerge. To foster the community interaction and sharing, we plan for additional dissemination in the future.

Board expectations

Boards typically care about:

- Strategic positioning and growth of the organization
- Shareholder value, brand protection
- Strategic plans, resource allocation, management compensation
- Oversight of compliance (government and sector regulations, ESG)
- Critical business risks - including cybersecurity
- Comparison with sector/peers
- Individual Board members' fiduciary liability.

For most of these areas there is an established practice of how to collect and report evidence in a way that is helpful, with an appropriate level of granularity and distribution of responsibility/delegation.

Regarding cybersecurity, the established practice in the industry is less mature. Often, Boards feel insufficiently competent to understand cyber risk or find cyber too technical, they approve resources and delegate this risk.

Boards often fail to see the continuous importance of cybersecurity and have knee jerk reactions to breaking cyber stories in the media then quickly forget about it until the next big cyber incident. Typically, cybersecurity only becomes an issue when it is already too late.

This is sometimes emphasized by an untransparent management culture, systematically reporting all clear / all green, whereas in reality most Boards would want to hear about gaps and how these gaps can be fixed.

In cases where cybersecurity reporting to the Board is taking place, there is a wide variety of methods, tools, and processes in use. Organizations struggle with what to report and how to obtain effective feedback from the Board.

It's all about risk

Our cyber environment requires us to make choices in terms of what to protect and how. Perfect security is an illusion and resources are scarce. Assessments and decisions regarding priorities are facilitated and objectivized by using the established practices of risk assessment.

There are different definitions of risk, centered around the potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its **likelihood** and the associated **impact**⁴. A real-life example of risk is the potential of dying or becoming seriously ill as the result of a pandemic virus infection.

For our discussion, we will use the expanded model in which risk is composed of three factors **Threat x Vulnerability x Impact**. In this equation, the likelihood is expanded into a combination of threat and vulnerability which in the cybersecurity context is helpful. We are not covering accidental events.

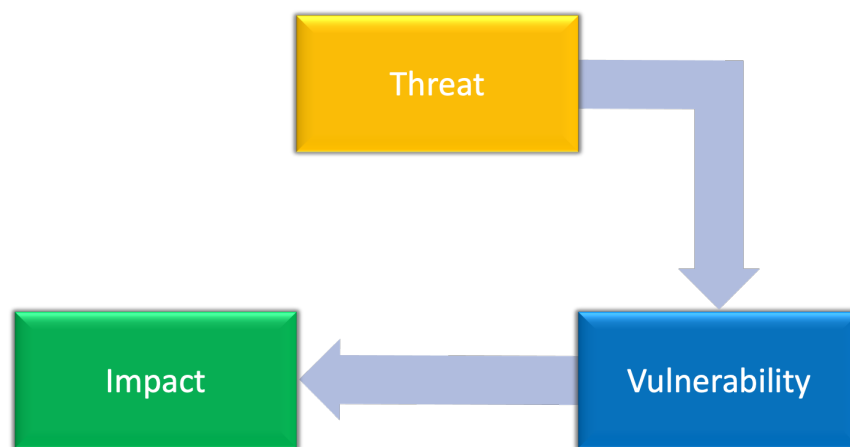


Figure 1 Risk as a combination of threat, vulnerability and impact

Threat is mostly external to our organization and is closely linked to adversaries. Identifying our key adversaries and their motives is important for prioritization. We can observe current threats and try to predict future threats. Specialized threat intelligence companies and government services can help to understand which adversaries we have, what their motives are, which tools and methods they use, and how they operate to pursue their goals.

The second factor is **Vulnerability**, and this is the one on which we can have the most leverage by designing and implementing controls. Identifying key controls,

⁴ <https://www.dhs.gov/xlibrary/assets/rma-risk-management-fundamentals.pdf>

considering our key assets, and the motivation and methods of our key adversaries is important for prioritization.

In terms of **Impact**, we can think of intellectual property theft, leakage of private data, interruption of service, personal harm, and brand damage. Impact is closely linked to assets. Identifying our key assets is important for prioritization.

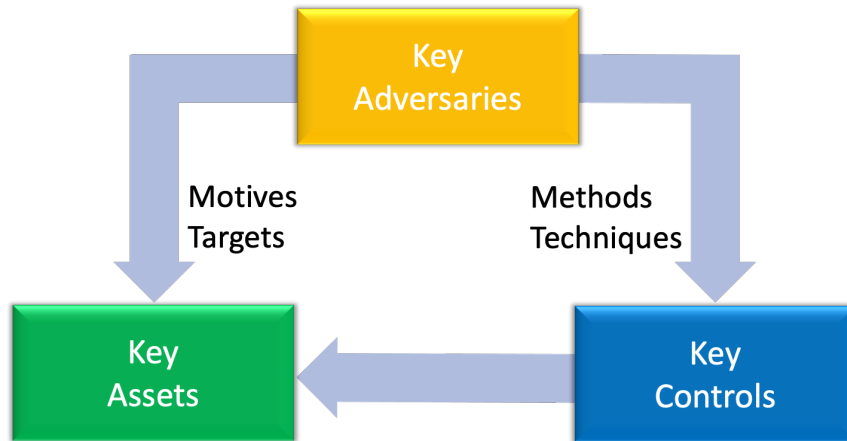


Figure 2 Cyber risk from adversaries targeting assets by exploiting vulnerabilities

We believe we should tackle the cybersecurity problem as a risk management problem and to use informed risk management and mitigation to prioritize action continuously. Cyber must be integrated in the overall management system, it should not be considered as something special/isolated but as an integral part of the organizational activities and processes including the risk management process. This requires alignment in methods and vocabulary.

To illustrate the relevant information flows in the current document, we will use the diagram in Figure 3, inspired by the NIST Cyber Security Framework⁵.

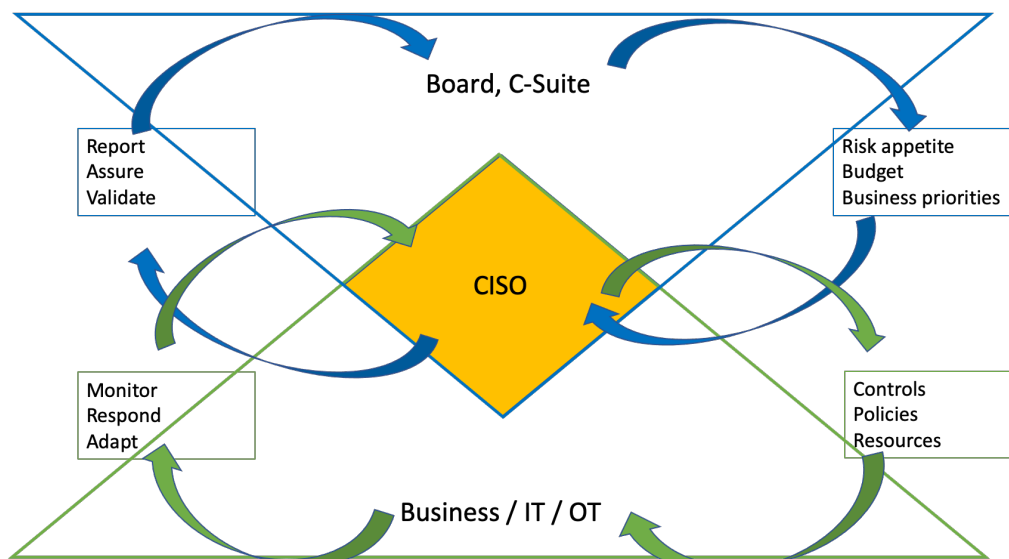


Figure 3 Information and Decision Flows. Inspired by NIST CSF

⁵ <https://www.nist.gov/cyberframework>

We can distinguish an upper, senior executive part and a lower implementation/operations part with the CISO in the central overlapping zone, connecting the cybersecurity operational level with the strategic level.

Important choices to make

It is necessary for organizations to make fundamental choices for cyber risk management. We can find these illustrated on the right-hand side of the diagram: which are the key assets, what is the risk appetite, and which are the key controls/mitigations to put in place. And related to those, the budget and the resource allocation to cybersecurity means and staffing.

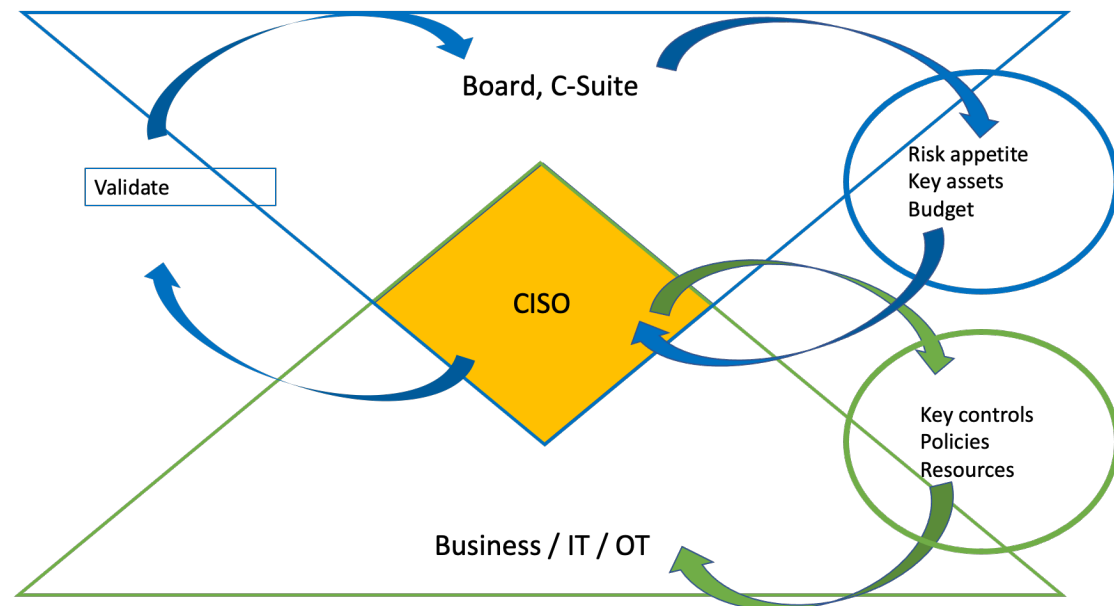


Figure 4 Choices to make

It is of crucial importance that these choices proposed by the CISO be agreed and aligned across the organization (cybersecurity, risk, IT/OT, business) as well as understood/approved and kept up-to-date at the executive level.

Key assets - crown jewels

In most organizations it is cost prohibitive to protect all assets against all possible cyber threats. Priorities need to be identified and resources allocated to the most relevant threats to the most important assets. Identifying these key assets is an essential component of business risk management in general and cybersecurity risk management in particular.

It is a non-trivial task activity, requiring cross-functional analysis and assessment, taking into account potential impact on business continuity, privacy, regulation and long-term competitive position (intellectual property).

When identifying (and updating) the list of key assets, a CISO should look beyond IT assets (data centers, backup systems, active directory etc.) and include relevant information assets (repositories, intellectual property), business assets (accounting, production management, logistics, physical access) etc.

Much is made of identifying the key assets or “crown jewels” which itself is a probabilistic risk assessment, an expression of the belief an attacker is more

likely to steal x than y. We say probabilistic here, because it assumes we can provide less defences to some assets than others. However, another key element of probabilistic reasoning is to update those assumptions based on current cyber threats. Crypto-jacking for example, does not care what your crown jewels are, and is quite content to just dwell on less important assets.

It is important to evaluate the probability of various cyber threats (DDoS, ransomware, targeted IP theft, opportunistic breach, phishing, fraud, malware infection – this list is non-exhaustive) in relation to discussing the key assets, one man’s trash is another man’s treasure.

Risk appetite

To identify risk mitigation and control measures, an organization needs to determine at the executive level what level of control is “good enough”, or what is an acceptable level of risk.

In doing so, we should express the risk appetite in a quantifiable way, using a threshold or a graphical representation of acceptable and non-acceptable situations. Some organizations use monetary thresholds for risk appetite. For other organizations (transport industry, hospitals, etc.), the threshold could be related to risk of injuries or loss of life. In some cases, the risk appetite may be related to business continuity or acceptable duration of interruption of service.

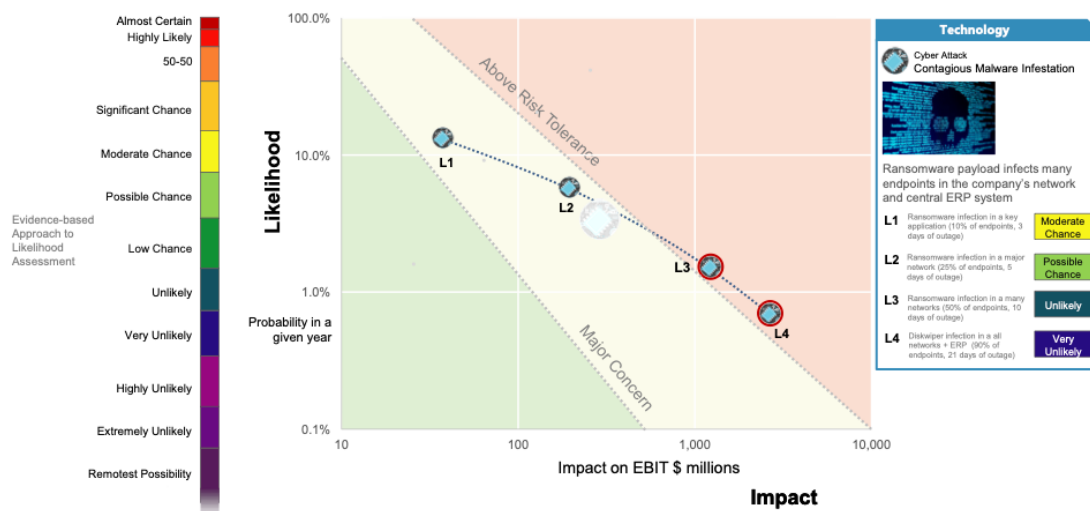


Figure 5 Example of mapping towards risk, credit Center for Risk Studies, University of Cambridge

Cybersecurity framework(s) and Key Controls

Cyber security frameworks are a tool to manage cybersecurity risks in a coherent manner and to implement a corporate cyber security strategy. Widely used control frameworks are ISO/IEC 27001⁶, NIST’s Cyber Security Framework (CSF)⁷, its derivative the CRI Profile⁸, NIST SP 800-53⁹ as well as the CIS Critical Security Controls¹⁰. Alternatively, or in combination, many organizations also use

⁶ <https://www.iso.org/isoiec-27001-information-security.html>
⁷ <https://www.nist.gov/cyberframework>
⁸ <https://cyberriskinstitute.org/>
⁹ <https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final>
¹⁰ <https://www.cisecurity.org/controls/cis-controls-list/>

the threat-centric MITRE ATT&CK[®] Framework¹¹. While it could be overwhelming to choose a single framework, they all have their specificities, it does not matter much which one is chosen because there are mappings between them. It is important though to choose one and to stick with it, so that the organization can measure progress over time.

It is highly advisable for any organization to seek internal agreement on the framework profile to use to frame its cybersecurity strategy and risk mitigation. Without such internal alignment between CISO, IT/OT, and risk management, it is difficult to engage the Board in cyber.

A good place to start when identifying and monitoring key controls is to map adherence to the baseline cybersecurity guidance of national cybersecurity authorities. We have included a selection of relevant sources in Annex 1. There is a large degree of overlap between these different sets of baseline guidance and they still need to be transposed to the specific situation of an organization. However, they do provide an excellent, succinct, and practical starting point.

Some key controls which are invariably included:

- K1: Maintain an up-to-date inventory on all (key) assets and dependencies;
- K2: Produce reliable, valid, safe, and secure backups of key assets;
- K3: Enforce multi-factor authentication wherever possible;
- K4: Limit users' access permissions to what is strictly necessary;
- K5: Identify and perform timely patching of important vulnerabilities;
- K6: Collect and analyze logs of all (key) assets;
- K7: Segment the network to protect key assets;
- K8: Harden internet facing systems;
- K9: Implement an incident response and recovery process;
- K10: Raise user awareness (including Board members).

We will refer to these key control identifiers in the metrics examples below.

Quantitative metrics

It makes sense to combine the selection of a framework profile with the definition of quantitative metrics (KPIs, KRIs, KCIs, OKR¹²) with goals/outcomes and link these to the relevant processes/systems and process owners. These metrics can be measured against the accepted goals over time, compared across the business, and benchmarked against peers. A few good industry practices show the potential of this approach:

- CIS Controls Measures and Metrics¹³;
- EPRI Cyber Security Metrics for the Electric Sector¹⁴;
- The Dutch payments association Library of Cyber Resilience Metrics¹⁵;

¹¹ <https://attack.mitre.org/>

¹² Objectives and Key Results, High Output Management, Andrew S. Grove.

¹³ <https://www.cisecurity.org/insights/white-papers/cis-controls-v7-measures-metrics>

¹⁴ <https://www.epri.com/research/products/000000003002010426>

¹⁵ <https://www.betalvereniging.nl/wp-content/uploads/Library-of-Cyber-Resilience-Metrics-Shared-Research-Program-Cybersecurity.pdf>

- The German automobile sector KPIs linked to ISO27001¹⁶;
- NIST's Performance Measurement Guide¹⁷.

Most frameworks assume that organizations implementing them use self-assessment, potentially combined with some form of external review by a certification or auditing body. Self-assessment is also in line with the standard practices in corporate risk management.

Monitoring by self-assessment has fundamental drawbacks to provide the status of the cyber risk mitigation measures and their effectiveness. These include:

- It is subjective (no separation of duty, same level of knowledge);
- It is not granular enough;
- It is time consuming;
- It is disconnected in time from the events;
- It cannot be used for alerting/escalation/response;
- It may require independent auditing to be acceptable by regulators;
- It may be limited to deployment indicators (what has been implemented?).

Machine-generated data can provide a very useful complement to self-assessment or even replace it to a large extent. They can make reporting on cybersecurity risk objective, repeatable, and automated. The identification of the machine-generated data sources and analytics needed for the metrics is an important step in the process of designing and implementing a coherent, comprehensive, and effective set of metrics.

The number one danger of metrics for cyber risk is that they begin to reflect work done or effort applied, instead of risk reduction. A Board or executive team must rigorously push back against the inclusion of such metrics. That is an operational matter, not a risk one. In other words, avoid metrics for number of incidents worked, or malware quarantined. These are fantastic operational metrics, but they do not tell the Board if the money they spent reducing the risk is effective. One easy aid is that a risk metric usually takes the form of a proportion or ratio, such as accidents per 1000 miles driven. If there is no ratio in a metric, then dig a little deeper on how it is measuring risk variance.

Also, sometimes it is good for a metric to show an increased risk. Early warning systems are a sign of a healthy risk team, and cyber risk is dynamic. Therefore, do not punish metrics or teams which communicate an increased risk, they may be carrying a very timely message to you.

A number of keywords come to mind when thinking of what makes good metrics:

- Objective
- Immutable
- Repeatable
- Continuous
- Relevant
- Effective

¹⁶ <https://www.vda.de/vda/en/News/publikationen/publication/vda-isa-catalogue-version-5.0.4>

¹⁷ <https://csrc.nist.gov/publications/detail/sp/800-55/rev-2/draft>

- Informed
- Agreed
- Actionable

A Metrics Model

We propose a Metrics Model with the following three steps:

1. Gather relevant cyber evidence
2. Transform the evidence into business risk¹⁸
3. Report to the Board, provide reasonable assurance, and highlight gaps.

In this Model, every step is deconstructed into building blocks which we will illustrate and comment below, and for which examples from the community are included in Annex 2. The aim is to provide inspiration and insights for organization-specific solutions rather than to infer that we propose a perfect solution to the cyber risk measurement and reporting challenge.

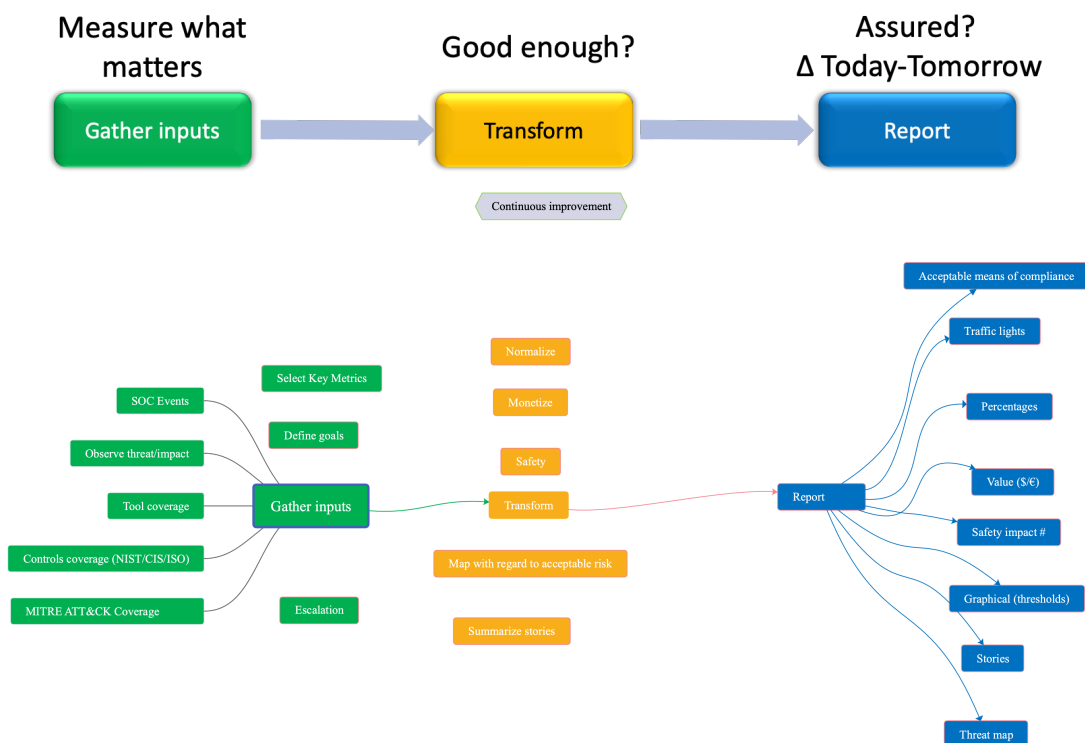


Figure 6 Metrics Model

Improvement loops (local and overall) should be incorporated in the Metrics Model and its relevant processes, adapting it to changes in stakeholder expectations as well as in the risk posture (threat landscape, vulnerabilities, dependencies). Insight and methods emerging from the community also lead to improvements.

¹⁸ Business risk is the exposure an organization has to factor that will lower its financial goals or lead it to fail. A business risk can be of many types such as strategic, operational, reputational, compliance, or financial.

Gathering inputs – measure what matters most

On the input side of the Metrics Model we find technical metrics, which should be (a subset of) the ones that are used by the business/operations to implement and monitor operational cyber risk mitigation. In Figure 7, we find these technical metrics on the lower left-hand side of the diagram.

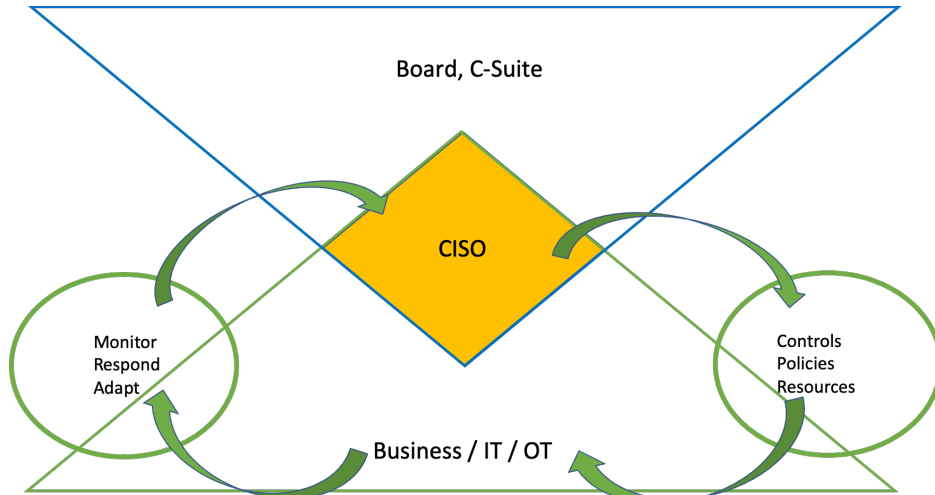


Figure 7 Technical metrics - inputs



Figure 8 Collecting inputs - building blocks

We can distinguish different families of operational key metrics which we group by nature (**control-centric**, **threat-centric**, **tool-centric**, and **event-centric**).

Control-centric

In this category, we find metrics which the organization identifies to measure alignment with a set of key controls. These are related to a control framework (NIST CSF, ISO, CIS, etc.).

Control-centric metrics could include:

- Coverage of a control - for all assets or a selected group of (key) assets;
- Effectiveness of a control;
- Data source and update frequency;
- Threshold level.

A (more granular) variation of this control-centric approach decomposes a control's coverage into three components: deployed, operational, and effective. It is important to note that these measurements should be continuous because the threat landscape and ability for the control to manage risk will shift over time. These three controls are defined as follows:

- Deployed – is the control installed where it should be;
- Operational – is the control functioning as designed;
- Effective – is the control is operating effectively, a measure ("evidence") of whether a given control is contributing to the reduction of the risk over a period of time.

On every one of these three domains, a score is established by collecting evidence. A combined score of the three areas gives a "Coverage Score".

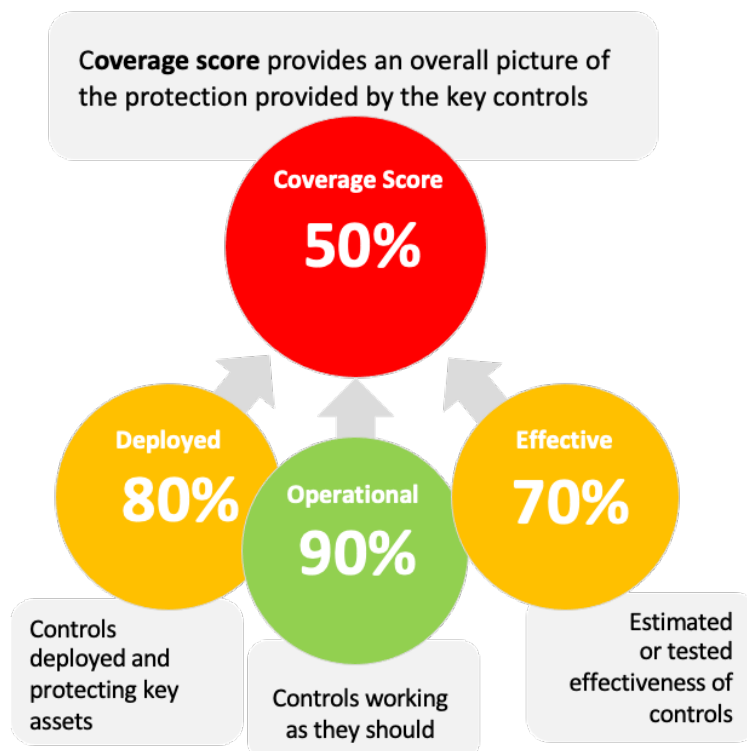


Figure 9 Example of a control-centric metrics in a pandemic infection

A real-life example of this concept can be found in pandemic infections, a vaccine is one of the possible key controls;

- Deployment would be the vaccinated part of the population (in this example 80%);
- If the vaccine generates an immune response only after a certain period, this generates a difference in the share of the deployed vaccines that are operational (in this case 90%);
- A vaccine is only effective to a certain degree (in this case 70%);
- Therefore, in this case the overall coverage is 50% (a combination of the three factors).

$$\eta = \frac{70}{100} \times \left(\frac{90}{100} \times 80 \right)$$
$$\eta = 50.4$$

The effectiveness of the controls could be tested on individual controls (pen testing) or on all controls deployed (Red Teaming). In the latter case, the result could be used to estimate the overall level of mitigation of cyber risk.

A control-centric approach will usually be found in heavily regulated environments. As an example, government departments in the US are expected to implement NIST 800-53 which includes some 1.000 controls and control enhancements.

However, even in regulated environments with mandatory controls, it makes sense to identify the key controls which matter most to the mitigation of the current cyber risk. Selection of key controls could be fostered by understanding the key threats (motives and techniques) and the key targeted assets.

Some examples of control-centric metrics

- K1: Percentage of (key) assets (endpoints, network, servers) inventoried;
- K1: Unencrypted databases storing personally identifiable information;
- K2: Percentage of key assets compliant with the backup policy;
- K4: Percentage of endpoints without local admin rights;
- K4: Percentage of endpoints with application white listing implemented;
- K4: Percentage of privileged accounts managed by an access control solution;
- K8: Percentage of internet-facing (key) assets scanned weekly for vulnerabilities and misconfigurations;
- K9: Percentage of critical applications without Business Impact Analysis;
- K10: Percentage of staff having followed cybersecurity training in the past year (including Board members);
- Effectiveness of key controls ascertained by Red Team or automated testing.

Threat-centric

In this category, we find metrics in which the organization identifies its most important adversaries and tracks the TTPs (Techniques, Tactics, Procedures) that they are known to deploy by using the MITRE ATT&CK® Framework.

Mitigation measures are mapped against these techniques in a similar fashion to the control-centric approach. This knowledge must be kept up-to-date with the latest information on notable adversaries and relevant incidents.

In the figure below the use of techniques by different relevant adversary groups is highlighted in color, from yellow (less prevalent) to red (techniques used by all relevant adversaries).

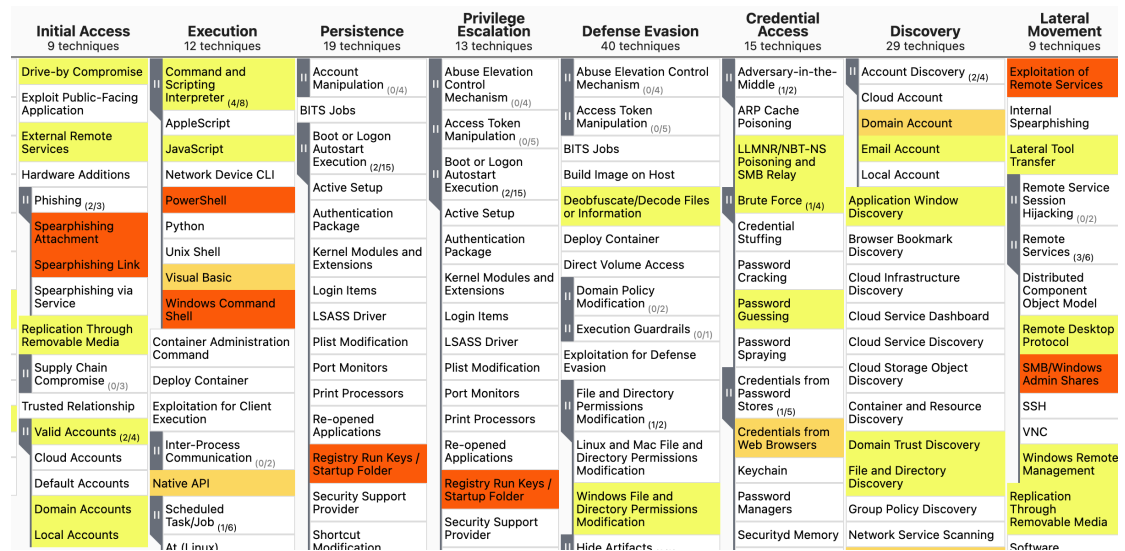


Figure 10 Example of the use of a heatmap to show prevalence of techniques

Threat-centric and control-centric approaches can be combined by using mappings between control and threat frameworks¹⁹ to identify select relevant threats and mitigating controls and convert them into key metrics.

Some examples of threat-centric metrics:

- Percentage of mitigation coverage of techniques known to be used by notable adversary groups;
- Percentage of coverage of key mitigations by an active testing program (automated or Red Team);
- Percentage of coverage of notable adversaries and their techniques with SOC playbooks and hunting programs.

Tool-centric

In this category, we find metrics in which the organization focuses on the deployment of specific cybersecurity tools (EDR, perimeter defenses, MFA, etc.) to achieve risk mitigation. The data collection on the deployment of tools is straightforward and the mapping of the effectiveness of every tool against known threats is also well documented.

Similar principles could be used as in the control-centric approach, using coverage/effectiveness or deployed/operational/actionable. A tool-centric approach could be a steppingstone to a more coherent and complete approach based on a framework (control-centric or threat-centric).

¹⁹ <https://github.com/center-for-threat-informed-defense/attack-control-framework-mappings>

Some examples of tool-centric metrics include the:

- K3: Percentage of implementation of multi-factor authentication (MFA);
- K6: Percentage of systems with full suite of security tools and policies (EDR, logging, gold standard software and configuration, policies, etc.);
- K6: Percentage of (key) assets with logging visibility;
- K8: Percentage of assets forced to connect to the internet via a proxy;
- Percentage of assets covered by automated controls and remediation.

Event-centric

Many organizations collect data on cybersecurity events (#alerts, #incidents, #false positives, #vulnerabilities, etc.). Such statistics can provide valuable input into the management of cybersecurity risk, but they need to be interpreted. Is it good or bad if more vulnerabilities are found or more incidents are happening? Have detection methods improved, or have the systems degraded?

Some examples of event-centric metrics include the:

- K1: Number of security systems implemented versus its asset coverage;
- K4: Number of issues found in monitoring/screening privileged assets;
- K5: Percentage of systems patched within SLA;
- K6: Number of false positives in the Security Operation Center;
- K8: Number of (external facing) orphaned assets found;
- K9: Number of critical incidents/average time to discover/contain;
- K9: Percentage of critical and high security alerts reviewed within SLA;
- K10: Number of corporate credentials in the wild (Account Take Over);
- Annual cost of cyber incidents;
- Number of open high-risk security and privacy issues beyond SLA without a remediation plan.

Timing data on incidents and vulnerabilities can provide useful information on the performance of the cybersecurity organization and systems. Good progress on this subject has been made in the First Metrics SIG²⁰.

Supply chain risk

More and more companies are experiencing the impact of cyber incidents affecting their suppliers, either directly through network connections or products, or indirectly through interruptions of the supply chain affecting business continuity. Mapping dependencies on suppliers, gaining insight into their cyber security posture, and implementing appropriate controls is becoming an integral part of cyber risk management and should therefore also be included in the metrics.

Monitoring the cyber risk of suppliers can be sourced from specialized companies and mitigations can, to a certain extent, take the shape of contractual terms and insurance coverage. However, a clear picture of dependencies and scenarios for detection and response should be established. Additional

²⁰ <https://www.first.org/global/sigs/metrics/events>

guidance can be found in the NIST publication on Key Practices in Cyber Supply Chain Risk Management²¹.

Some examples of supply chain metrics include the:

- Percentage of critical vendors/suppliers for which inventory of assets, dependency, risk assessment, and mitigation has been performed;
- Percentage of critical vendors/suppliers with security annexes;
- Percentage of critical vendors that have been audited;
- Number of critical vendors/suppliers with open high-risk security and privacy audit findings without a documented risk management plan.

Observe impact - Stories

Many organizations document relevant incidents (internal, impacting peers, sector, or region) using narratives in “stories”. This kind of anecdotal evidence is very appealing for non-technical C-Suite and Board members because they exemplify what may happen (or has happened) to the organization. They also allow the CISO to draw the attention to trends in frequency, impact, methods in the threat landscape, and support prioritization of action in terms of controls and resource allocation.

Select Key Metrics and goals carefully

One must choose key metrics carefully because the selection and reporting of metrics drives an organization. The selection of metrics infers what matters most to the leadership and people will align to them. Measuring the wrong things will in turn drive against the desired cybersecurity goals and lead to false assumptions about the risk posture. In addition, the Board may want to focus its attention on improving the indicators rather than the underlying cyber risk posture.

Key metrics should evolve over time with the increasing maturity of the organization, changes in regulatory requirement, business objectives, and changes in the cyber threat landscape. For the selected metrics, goals must be set and agreed across the organization. These need to make sense in term of risk mitigation and risk appetite. They could include a timing component in case the organization wants to include an evolution in maturity over time.

Escalation process

It is recommended to define a process/threshold which triggers emergency reporting of deviations/developments to the executive level in between reporting periods. Obviously, one can think of critical incidents/breaches, but the trigger could also come from important vulnerabilities or developments in the threat landscape requiring immediate attention from the executive level. A recent example of such a case was the Log4j vulnerability and many organization’s dependencies on products using this software component.

An escalation, process could also be designed for metrics which are only reported at Board level in case a predefined threshold is breached. This could decrease the overload of irrelevant information delivered to the Board.

²¹ <https://nvlpubs.nist.gov/nistpubs/ir/2021/NIST.IR.8276.pdf>

Data sources – undeniable truth

Data collected from inside

Technical metrics should be composed of data that is automatically collected from the source infrastructure, with minimal human involvement. These include:

- Asset management and discovery systems (completeness, criticality);
- Tool management systems and consoles (deployment);
- Logs and SIEMs (deployment and operation);
- Scanning software (versions, vulnerabilities, configurations, policies);
- Identity, privilege, and access management (controls and policies);
- Network traces (completeness, controls).

Most of this data pertains to the deployment of controls, tools, and policies. As for the effectiveness of the risk mitigation they can be derived in a theoretical manner on the basis of expected mitigation from a specific control.

Data collected from testing

Additional insight on implementation (deployment and operation), in particular its effectiveness, can be gained from testing the controls. Typically, such testing of controls can be done by human pen testing/Red Teaming or automated testing using specific tools or bug bounty programs. This category of metrics will particularly make sense in an organization which already has a mature Information Security Management System²² in place.

Data collected from outside the infrastructure

Some data on confirmed infections and vulnerabilities can be collected from outside an organization's infrastructure by scanning or observing network traces calling out to known malicious infrastructure.

Transforming – Are our controls good enough?

Whereas operational key metrics are important for the CISO to steer the implementation of the cybersecurity strategy and to monitor controls in a granular manner across the board, these are not appropriate for reporting to the executive leadership, Board, and other strategic stakeholders. They would be perceived as overwhelming, cryptic, and disconnected from the business risk.

In order for cyber risk metrics to resonate at the Board level, they need to be transformed to meaningful business reporting (money, safety, brand value, etc.) and compared with the risk appetite. Is our risk mitigation good enough? Can we provide reasonable assurance? Can the Board validate our assumptions and orientations? In the following figure we show the flow of information from technical/operational metrics towards Board metrics.

²² <https://www.iso.org/isoiec-27001-information-security.html>

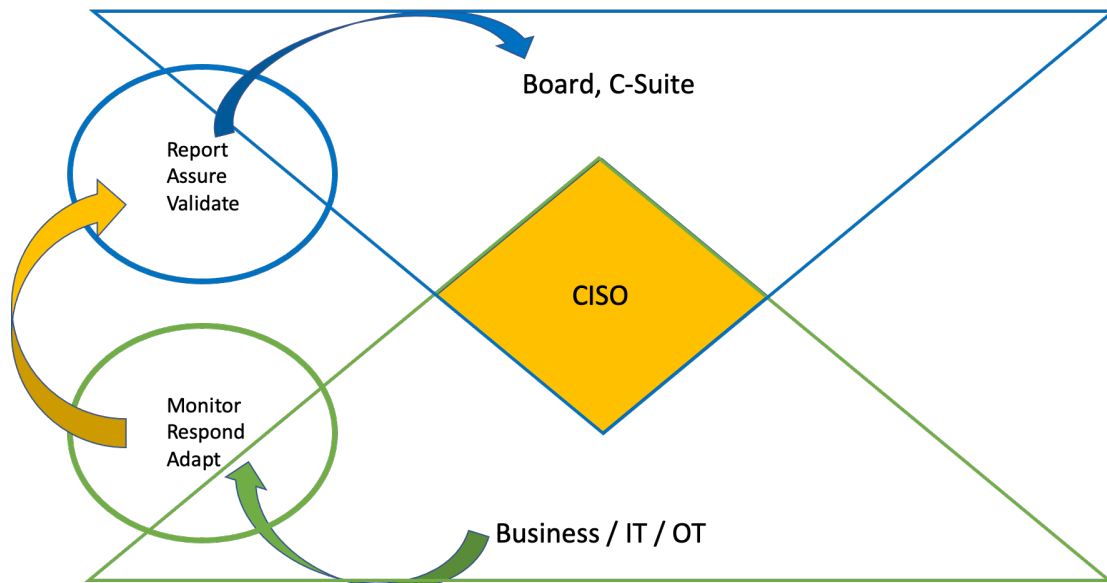


Figure 11 Transforming technical metrics into strategic metrics

Again, we distinguish a number of building blocks in the Transform step. These convert operational key metrics into values that can be compared with tolerable risk, be integrated into business risk, and be reported to the Board.

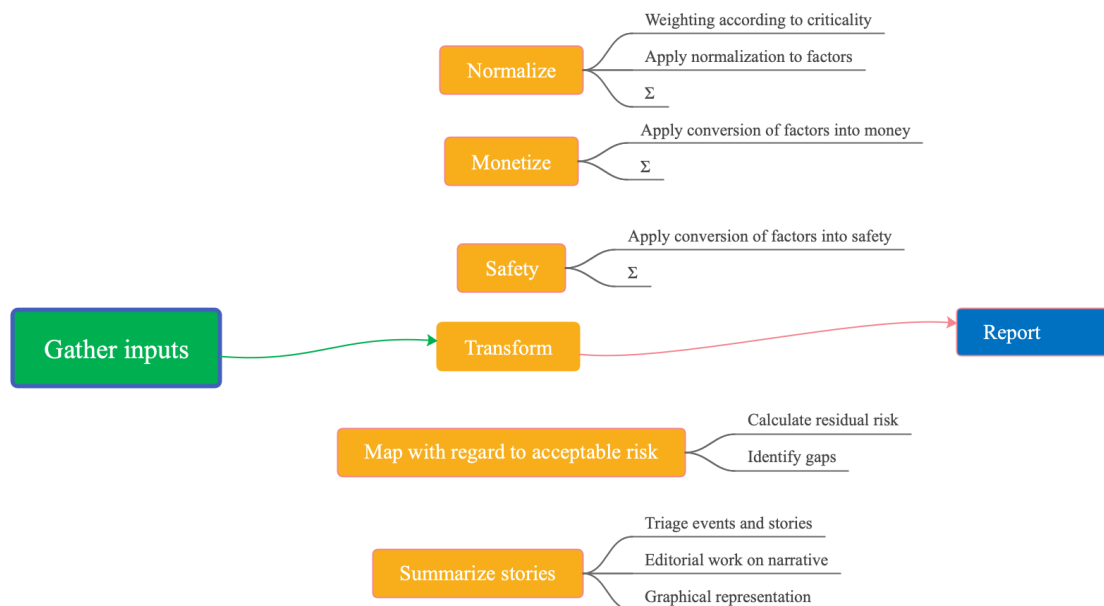


Figure 12 Transform - building blocks

Normalize

In any organization there might be a large number of metrics to describe the state of controls and organization performance. This can cause a negative effect where stakeholders become overwhelmed by detail. Another issue is that the value of one measurement alone can be meaningless when viewed independently, but becomes critical when viewed across a set of metrics.

For example, if you want to understand the health of your endpoints against malware, viewing the deployment of antivirus software on a given operating system alone will not be enough. You would need to view the different measurements across the different operating systems. Additionally, anti-virus

software alone will not provide the answer. You will need to view measurements from other tooling such as Event Detection Response (EDR).

To overcome this challenge, sets of metrics, even though they are different in nature, can be normalized or harmonized to provide a more holistic view. Such normalization should provide a simplified view on a large number of distinct control domains while also revealing insights in important gaps which could become obfuscated by consolidation.

Normalization could also include a weighting component to take into account different levels of criticality of assets. For example, the coverage of controls in highly critical assets could be assessed as more important than in other assets. In a consolidated metric this could be accounted for by weighting.

Using a three-tier (red/amber/green) min/max scaling normalization, it is possible to remap the metric to a new scale, while maintaining the exact proportion within each tier (i.e. an input metric that is at the top end of red with remain in the top end of red even though its numeric value may change).

In this graphic several low-level metrics are normalized to a common scale. Once normalized to a common scale, these metrics can be meaningfully aggregated or combined and these aggregations can be cascaded to obtain just a few top-level summarized metrics.




Figure 13 Three-tier (red/amber/green) min/max scaling normalization

Monetize (value at risk)

In organizations, business risk appetite is expressed in monetary terms. As such, it is necessary to attempt to transform metrics to enable digestion by key stakeholders to be able to work within normal business processes. This is particularly impactful when asking for funding. For example, if cybersecurity needs \$10m to deploy a new EDR system, identifying that it will reduce the current risk to the organization by \$12m and place the company back within risk appetite is a strong justification.

This step converts individual or aggregated key operational metrics into value of risk managed. In terms of value, we need to consider the direct impact of the

incident on the business (continuity of operations, litigation by customers, compliance penalties, cost of incident response, ransomware payments) as well as indirect impact (brand image damage, stock price).

One of the methods to perform monetization is Factor Analysis of Information Risk (FAIR™)²³, a model for understanding, analyzing, and quantifying cyber risk and operational risk in financial terms. This method is well established and thoroughly publicized. However, it may be too elaborate and difficult to maintain for smaller or less mature organizations.

The Dutch telecommunications provider KPN's Potential Harm of Security Incident (PHOSI) calculator is a low entry alternative, available as an app on smartphones^{24,25}. It facilitates the calculation of value at risk from a small set of questions. These PHOSI estimations can be combined with individual threats/controls or also with results from Red Teaming (which potential harm was avoided by timely patching a critical vulnerability or by performing a Red Team exercise?) or exposure from important vulnerabilities.

A broader and simpler way to monetize key operational metrics is to use the result of the normalization/consolidation as the mitigation factor to be multiplied with the average frequency and impact of a cyber incident observed in the community. Research work on the frequency and impact of ransomware incidents has resulted in interesting examples for such approximations²⁶. Recent academic work in this domain can also be found in "A System to Calculate Cyber-Value-at-Risk"²⁷.

It is important to note that this is an active area of research and the methods to quantify risk are imperfect. While this is a goal to strive for, and essential for communicating with key stakeholders, the results should be treated with care.

Safety impact (life at risk)

Some organizations are, because of their activity, not only concerned with monetary impact but also with safety impact. This would be the case for airlines/air traffic management, car manufacturers, hospitals, nuclear energy suppliers, etc.

Cyber risks which could lead to loss of life would deserve to be estimated and the effectiveness of controls and risk mitigation measured and controlled. Much less work has been published in this domain, but the underlying principle would be similar to the value at risk calculation.

This is certainly an area in which organizations would feel less comfortable to share assessments or expose compromises and calculated risks. An external perception of acceptance of risks of loss of life by an organization could very quickly lead to brand image damage.

Monetization of loss of life not being acceptable (at least in some regions of the world), the notion of acceptable risk of losing life is assessed by a mix of

²³ <https://www.fairinstitute.org/what-is-fair>

²⁴ <https://apps.apple.com/us/app/kpn-ciso/id112223795>

²⁵ <https://play.google.com/store/apps/details?id=com.kpn.ksp&hl=en&gl=US>

²⁶ <https://www.youtube.com/watch?v=kSi-oXg4xV0>

²⁷ <https://www.sciencedirect.com/science/article/pii/S0167404821003692>

quantitative and qualitative means to calculate risk with the aim to reach no loss as far as reasonably practicable and tolerated by regulations.

Map towards risk appetite

The outcome of the value/life at risk assessments needs to be compared with the risk appetite of the organization. In many organizations this risk appetite has already been established within the business risk processes. If this were not the case, the CISO should prompt the business/Board to determine the risk appetite:

- How much are we willing to lose in the event the risk materializes?
- To what extent do we want the risk to be mitigated?
- How many resources are we willing to make available for mitigation?
- Do we insure part of the risk?

Some will probably argue it is not possible to assess the probability of a breach. However, we must remember that if there is no benchmark of risk appetite, no one will try to quantify the probability at all. We advocate initiating a discussion about risk tolerance such as “less than a 5% chance per annum of a cyber breach loss exceeding 1 million dollars”. While acknowledging that this kind of quantitative approach to cyber risk is difficult to accomplish and is the exception rather than the norm, the goal is aspirational both to engender risk quantification, but also to allow reasonable budgets.

Be wrong about these numbers first, and let the executives work toward answering the question in a repeatable methodology. If you have no idea what numbers to use, examine some of the other risks within your organization such as the risk tolerances for fire, flood, or workplace accidents. They may be very different risks, but they can give you a guide to how to express the risk tolerance you hope to achieve. You may discover in the end that your organization really has closer to a 10% chance, but then the discussion of how much it costs to reduce it becomes tractable and rational.

The mapping of value/life at risk towards risk appetite is necessarily a multi-dimensional analysis in which expected frequency and possible impact are combined and in which the current situation could be expressed in a number of data points based on hypotheses about the effectiveness of the controls.

The mapping can also be used to show paths of improvements related to proposed controls/investments. None of that is likely to happen though, unless the Board sets a risk tolerance in the first place. Only then can discussions be had about the realism of the numbers and expectations of cyber risk resilience.

Summarize stories

Selecting relevant stories (incidents inside and outside the organization, cyber threat intelligence, regulatory developments) and extracting the essence (why is this relevant?) is an essential complement to the quantitative metrics.

Narrowing down the narrative and projecting a convincing story requires specific skills from the CISO and his/her team. Selected stories need to provide additional context to the risk posture of the organization and serve a purpose. Otherwise, they risk to divert attention and capture energy and resources that could be better used elsewhere.

Reporting cyber risk –provide reasonable assurance

Reporting cyber to the Board should serve the purpose of (re)assuring the Board that the cyber risk is within the risk appetite today and tomorrow:

- Are we good enough?
- Are the resources allocated to cyber appropriate and effective?
- How do we compare with our peers and our sector?

Bearing in mind that the Board and its committees are not specialized in cyber, it would be advisable to help the Board to ask the right questions and do not overwhelm them with information. To compare the situation with images, the next image shows a “CISO” cockpit with operational instruments and consoles which allow the pilots to interact with the plane to bring the passengers safely and timely to their destination.



Figure 14 Credit Aeropers / Pilots of Swiss

The following image shows a “Board” cockpit with different instruments and without the possibility for control, being entirely dependent on ground control.



Figure 15 Credit NASA/SpaceX

The Board would expect the CISO to signal any developments that would substantially change the situation for better or for worse and to propose relevant actions and resources as a consequence.

Establishing a package of cyber metrics and reporting its context consistently to the Board, its individual members, and its relevant committees (Audit, Compliance) can be an effective and reliable way of providing cyber assurance.

It can be combined/aligned with reporting on other kinds of business risk as well as reporting on digital transformation strategy.

Metrics and narrative

The old adage “a picture is worth a thousand words” is also true in the context of cyber engaging your Board. A large diversity of graphical representation of cyber metrics is in use in the community and looking at such examples and interacting with peers in the industry can be very inspiring in the design of an organization’s reporting package. In Figure 16 we have included a few building blocks, based on examples from the community as included in the Annexes.

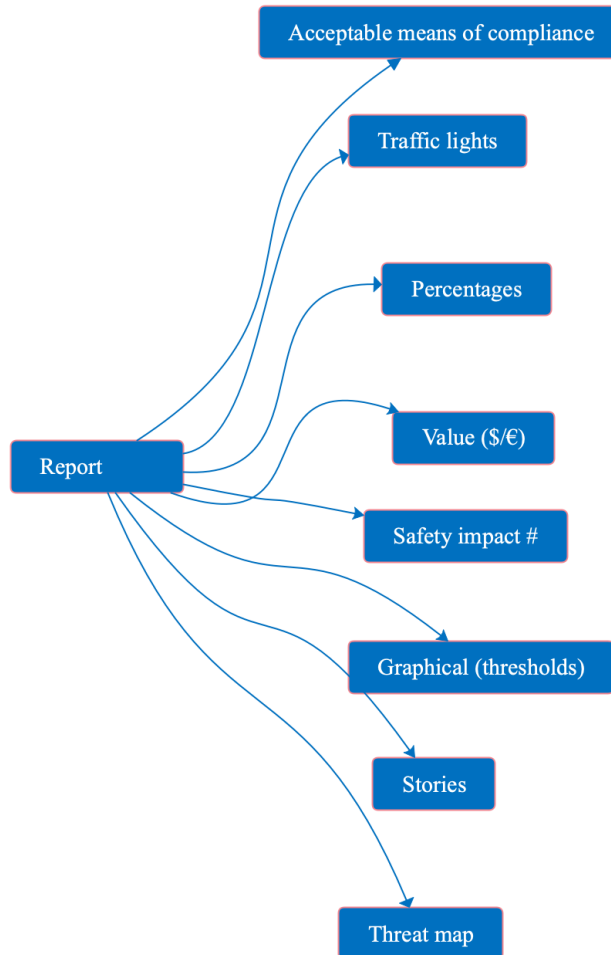


Figure 16 Reporting - building blocks

Communication channel(s)

Ideally an organization would establish a coherent and integrated communication channel to gather inputs, transform and then report them to the Board. The communication flows in Figure 17 would be implemented as intended. This requires different functions to cooperate and align on the framework, the processes to populate the metrics, and the roles and responsibilities in reporting.

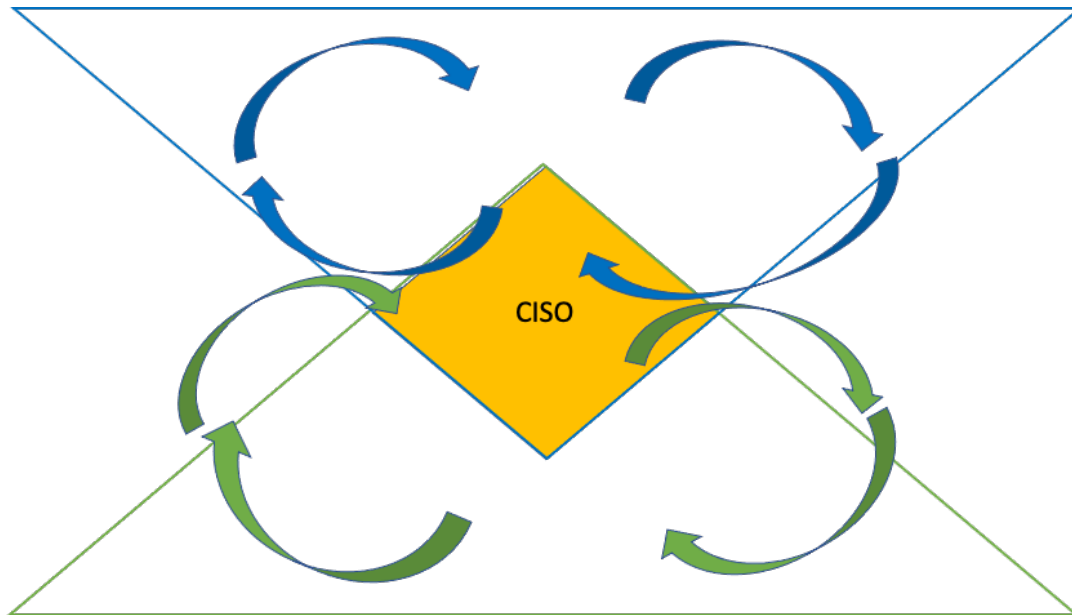


Figure 17 Optimal communication flow

Regardless of who actually reports to the Board, the CISO needs to play a key role in the process, assuring a professional and independent view on the cyber risk. In the optimal case it would be the CISO who actually reports to the Board in person, – allowing for interaction and personifying assurance.

Alternative situations, as illustrated in Figure 18, in which in the first diagram no information is provided to the Board, in the diagram in the middle information provided to the Board is not grounded in reality, or in the third diagram contradictory information is reported to the Board through different channels, are to be avoided or phased out.




Figure 18 Inexistent, unconnected, parallel communication flows

Overcoming resistance

While the concepts described in the current document are to a certain extent already used by cybersecurity departments, they are also a basis to integrate cybersecurity risks as part of the business risk management processes. However, these interactions are still challenging in many organizations where:

- Cybersecurity and risk are managed by different departments/silos;
- Cybersecurity is considered as a discipline for insiders, a “black art”, by the business departments;
- Risk management models are considered as complex and abstract by the cybersecurity departments;
- Cyber vocabulary and metrics are not translated into business terms.

Landing cybersecurity as a recurring issue on the Board agenda requires effort and convincing, but it can be facilitated by a number of proactive initiatives:

- Understanding the expectations of the Board and its individual members and adapting to changing expectations;
- Awareness raising sessions with the Board, explaining threats and risks in an understandable way;
- Incident response exercises involving the Board;
- Sending the Board monthly cyber briefs with relevant stories and context;
- Bilateral briefings with individual Board members expressing interest;
- Starting gradually and improving the system over time (for example, starting with the implementation of a maturity model-based approach before implementing a full-blown quantitative model);
- Transparent and positive cooperation with the Audit Committee.

Allocating resources to metrics and reporting

There is no doubt that implementing and maintaining a metrics and reporting system as described in this document requires dedicated resources. However, it also saves resources by internal alignment/streamlining, by avoiding needless reactive media responses and ultimately by focusing resources on what really matters to reduce cyber risk to an acceptable level, avoiding incident response and negative impact.

Annex 1: Where to start?

Questions for the Board to ask

- Do we have an inventory of key assets?
- Who is targeting us (key adversaries) and why?
- Which are our key controls and what is their status?
- Where are the gaps and how do we plan to close them?
- Do we have an incident response / business continuity / resilience plan?
- How much is at risk?
- How do we compare with our peers?

Key control baselines

Below is a (non-exhaustive) selection of baseline recommendations:

- Top seven security measures (Cybersecurity Centre Belgium)²⁸
- Top ten (UK National Cyber Security Centre)²⁹
- Essential eight (Australian Cyber Security Centre)³⁰
- Top 42 measures for a healthy network (FR ANSSI)³¹
- Top eight security measures (NL National Cyber Security Centre)³²

These authorities are also a source for up-to-date information on the threat landscape and the evolving nature of vulnerabilities and adversarial techniques.

²⁸ <https://cyberguide.ccb.belgium.be/en/take-security-measures-0>

²⁹ <https://www.ncsc.gov.uk/files/2021-10-steps-to-cyber-security-infographic.pdf>

³⁰ <https://www.cyber.gov.au/acsc/view-all-content/essential-eight>

³¹ <https://www.ssi.gouv.fr/en/guide/40-essential-measures-for-a-healthy-network/>

³² <https://www.ncsc.nl/onderwerpen/basismaatregelen>

Annex 2: Examples from the community

Gathering Inputs

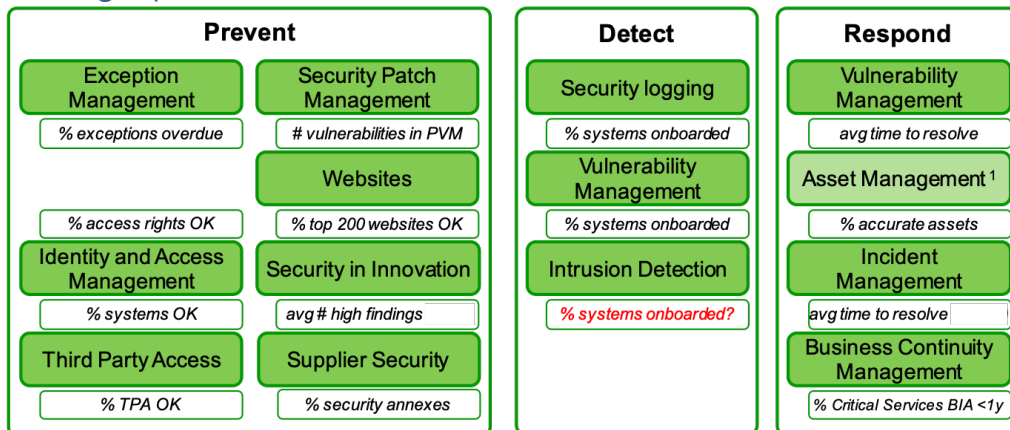


Figure 19 Example of event-centric metrics

Incident Timeline	Applicability Level	Description
Time of First Activity	Recommended for significant incidents	This is the earliest event in a confirmed or potential chain of events, that caused the incident.
Time of Detection	All incidents	The time that a control (e.g. telemetry, technology) or another detection mechanism (e.g. a human) recognizes that something has occurred.
Time of Containment	All incidents that require Containment	Time of Containment is the point in time at which the incident can no longer spread nor do damage.
Time of Remediation	All incidents that require Remediation	Time of Remediation is the point in time at which an affected target asset is returned to its pre-incident state or removed from the environment permanently.

Figure 20 Recommended timing metrics. Source FIRST Metrics SIG.

Function	Category	Subcategory	Subcategory Risk	Sub Cat Score	Subcategory Composite Risk Score
	Asset Management (ID.AM): The data, personnel, devices, systems, and facilities that enable the organization to achieve business purposes are identified and managed consistent with their relative importance to business objectives and the organization's risk strategy.	ID.AM-1: Physical devices and systems within the organization are inventoried	Critical	6	30
		ID.AM-2: Software platforms and applications within the organization are inventoried	Weighted Medium	4	
		ID.AM-3: Organizational communication and data flows are mapped	Weighted Medium	4	
		ID.AM-4: External information systems are catalogued	Critical	6	
		ID.AM-5: Resources (e.g., hardware, devices, data, and software) are prioritized based on their classification, criticality, and business value	Critical	6	
		ID.AM-6: Cybersecurity roles and responsibilities for the entire workforce and third-party stakeholders (e.g., suppliers, customers, partners) are established	Weighted Medium	4	
	Business Environment (ID.BE): The organization's mission, objectives, stakeholders, and activities are	ID.BE-1: The organization's role in the supply chain is identified and communicated	Weighted Low	2	10
ID.BE-2: The organization's place in critical infrastructure and its industry sector is identified and communicated		Low	1		
ID.BE-3: Priorities for organizational mission objectives					

Figure 21 An example of implementation of metrics linked to NIST CSF

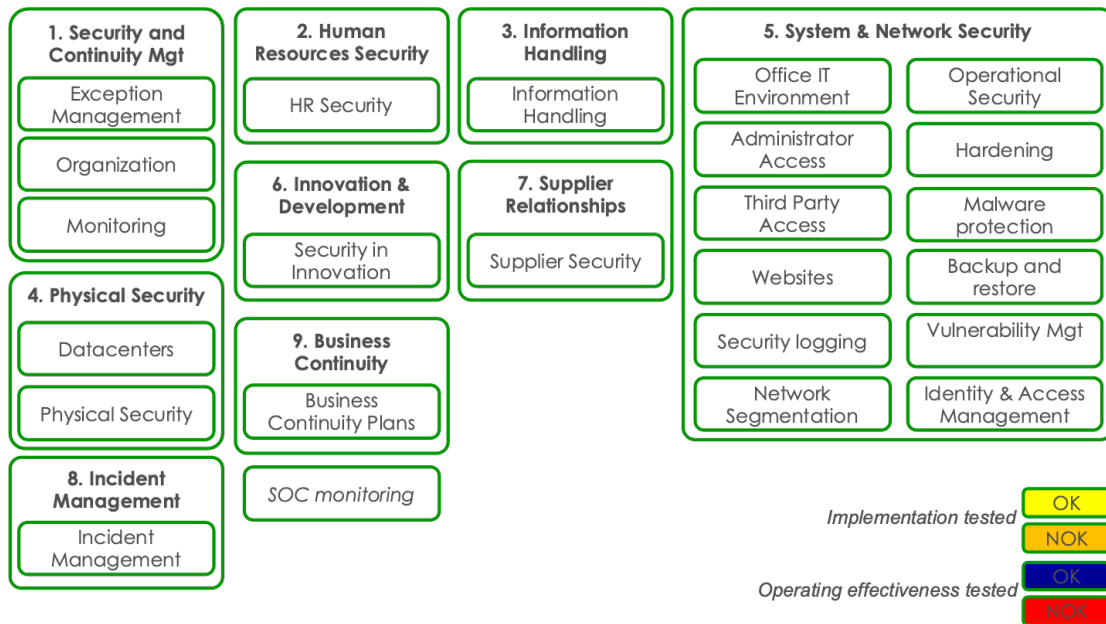


Figure 22 Example of metrics from testing

Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movement	Collection	Command and Control	Exfiltration	Impact
Spearpishing Attachment	Command-Line Interface	Registry Run Keys / Startup Folder	Scheduled Task	Obfuscated Files or Information	Credential Dumping	System Network Configuration Discovery	Remote Desktop Protocol	Input Capture	Remote File Copy	Data Compressed	Data Encrypted for Impact
Valid Accounts	Scripting	Scheduled Task	Valid Accounts	Scripting	Input Capture	Process Discovery	Remote File Copy	Data from Local System	Commonly Used Port	Data Encrypted	Disk Structure Wipe
Drive-by Compromise	PowerShell	Valid Accounts	Process Injection	Valid Accounts	Brute Force	Account Discovery	Pass the Ticket	Data Staged	Standard Application Layer Protocol	Data Transfer Size Limits	Resource Hijacking
External Remote Services	Scheduled Task	New Service	New Service	Code Signing	Credentials in Files	File and Directory Discovery	Remote Services	Email Collection	Connection Proxy	Exfiltration Over Command and Control Channel	System Shutdown/Reboot
Spearpishing Link	Exploitation for Client Execution	External Remote Services	Accessibility Features	Deobfuscate/Decode Files or Information	Credentials from Web Browsers	Network Service Scanning	Windows Admin Shares	Audio Capture	Web Service	Exfiltration Over Alternative Protocol	
Exploit Public-Facing Application	User Execution	Create Account	Bypass User Account Control	File Deletion	Network Sniffing	Remote System Discovery	Windows Remote Management	Automated Collection	Custom Command and Control Protocol		
Supply Chain Compromise	Windows Management Instrumentation	Redundant Access	Web Shell	Masquerading	Account Manipulation	System Information Discovery	Component Object Model and Distributed COM	Data from Information Repositories	Multi-Stage Channels		
Trusted Relationship	Dynamic Data Exchange	Web Shell	Exploitation for Privilege Escalation	Process Injection		System Network Connections Discovery	Exploitation of Remote Services	Video Capture	Standard Non-Application Layer Protocol		
	Rundll32	Accessibility Features	DLL Search Order Hijacking	Connection Proxy		System Owner/User Discovery	Pass the Hash	Screen Capture	Uncommonly Used Port		
	Service Execution	Bootkit	Application Shimming	Redundant Access		Network Share Discovery		Data from Network Shared Drive	Fallback Channels		
	Graphical User Interface	Component Firmware		Rundll32		Permission Groups Discovery			Multi-hop Proxy		
	Mhta	BITS Jobs		Software Packing		Security Software Discovery			Data Obfuscation		
	Regsvr32	Modify Existing Service		Web Service		System Service Discovery			Domain Fronting		
	Execution through API	DLL Search Order Hijacking		Bypass User Account Control		Virtualization/Sandbox Evasion			Data Encoding		
	Component Object Model and Distributed COM	Shortcut Modification		DLL Side-Loading		Query Registry			Domain Generation Algorithms		
	Windows Remote Management	Windows Management Instrumentation Event Subscription		DLL Search Order Hijacking		Network Sniffing			Standard Cryptographic Protocol		
	CMSTP	Winlogon Helper DLL		Hidden Files and Directories		Peripheral Device Discovery					
	Compiled HTML File	Account Manipulation		Hidden Window							
		Application Shimming		Indicator Removal from Tools							
		Hidden Files and Directories		Indicator Removal on Host							
				Modify Registry							
				Mhta							
				Network Share							
				Connection Removal							
				Process Hollowing							
				Regsvr32							
				Rootkit							
				Template Injection							
				Virtualization/Sandbox Evasion							
				Binary Padding							
				BITS Jobs							
				Disabling Security Tools							
				Execution Guardrails							
				Compiled HTML File							
				Component Firmware							
				CMSTP							
				Clear Command History							
				Compile After Delivery							

Figure 23 Example of the use of heatmap to show coverage of TTPs

Transform

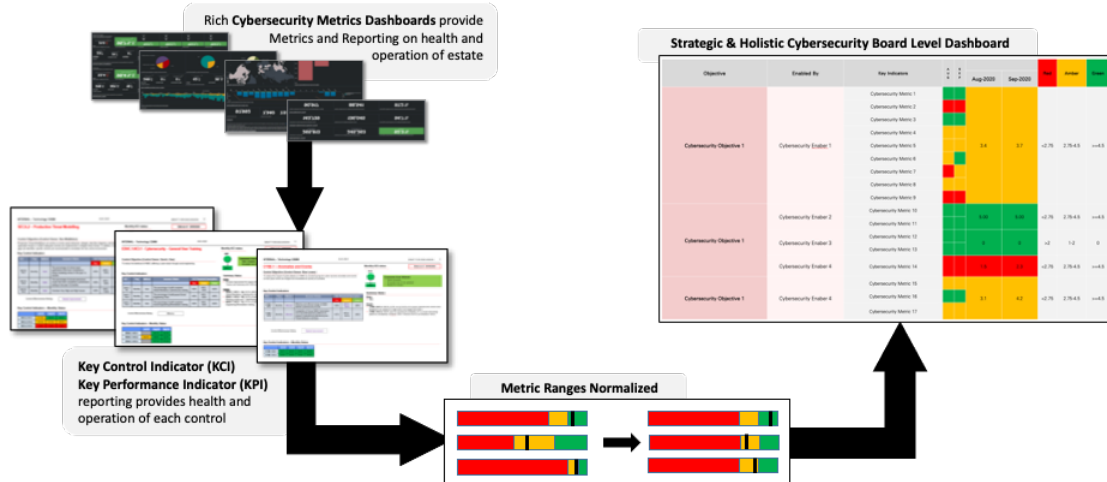


Figure 24 Example of normalization

Calculate **Potential Harm Of Security Incidents** for each security incident by two factors:

Likelihood:
How often would this vulnerability be exploited?

Potential Loss:
What would each exploit cost the company?

PHOSI = Likelihood x Potential Loss

Figure 25 Example of monetization





Figure 26 Example of mapping towards risk

Reporting cyber risk



Manufacturing industry or geographically close

When?	What?	Category	Group
Q1.2021	Hackers exploit IT tool Z to establish persistence	1	Unknown
Q4 2020	Ransomware attack at enterprise X	2 6	Cyber-Crime
Q4 2020	Enterprise Y hit by ransomware, data leaked	2 6	Cyber-Crime

Figure 27 Example of threat landscape reporting

When	Type	What?	Status	Group
4/2020	3	Spear phishing campaign with malicious Excel attachment.	Closed	APT X

Figure 28 Example of incident reporting

Group	Motive	Trend
Adversary 1	Adversary known to steal intellectual property in high tech industry.	↗
Adversary 2	Adversary known to steal intellectual property in our sector.	→
Targeted Cyber-Crime	Ransomware actor increasingly prevalent and sophisticated	↑

Figure 29 Example of notable adversary tracking

Threat Landscape Report 2021 Q3 – Executive Summary Direct Threats to EU Institutions, Bodies, and Agencies

INCIDENTS

4 significant incidents affected EUIBAs this quarter. In 3 cases the attack started with a compromise of a publicly accessible server (Oracle WebLogic, Microsoft Exchange).

In the other case, attackers obtained credentials via a phishing campaign.

In at least 3 significant incidents, threat actors successfully exfiltrated data.

Since the beginning of 2021, CERT-EU has already recorded 15 significant incidents, compared to 13 during the whole of 2020 and 8 in 2019.



THREATS

CERT-EU released 26 threat alerts (compared to 20 during Q1 and 22 in Q2).

The top 5 reasons for threat alerts were:

- Active exploitation of zero-days or n-days: Microsoft Exchange, VPNs, etc.
- Recent activity or new tools used by top threat actors
- Sharing actionable data related to TTPs used in significant incidents
- Spear-phishing campaigns directly affecting EUIBAs or sectors of interest
- Active use of commercial mobile spyware




Top threat actors: CERT-EU currently tracks 13 top threat actors. The level of exposure of EUIBAs has been high for 4 of them: two alleged Russian threat actors, one alleged Chinese, and one allegedly of North Korean origin.

Social media: The most frequently used social media network for impersonation of EU staff or the digital identities of EUIBAs has been Instagram, followed very closely by Facebook and at a distance by Twitter.

Malware and tools: The three most observed pieces of malware or malicious tools to which EUIBAs have been exposed were Cobalt Strike, Mimikatz, and Dridex. However, no infections have been confirmed.

Threats in Europe



Ransomware

A supply chain attack conducted by REvil against Kaseya VSA, software used by many MSPs, had a major impact on several organisations including in Europe, causing significant disruptions.

Taking into consideration the first 9 months of 2021, the average number of ransomware victims per month increased by 129% in 2021, compared to last year.

Nation-state activity

The EU has acknowledged and condemned Russian "Ghostwriter" cyberespionage / information operation activity against EU member states. The Russian APT29 threat actor targeted European governments with a zero-day exploit earlier in 2021. The EU, the UK, and the US attributed the Hafnium ProxyLogon attacks to China and are calling for an immediate stop to such adversarial activities. France reported a significant cyberespionage campaign by the Chinese Zirconium (aka APT31) threat actor. The NSD group and its Pegasus spyware have been used in several espionage cases against politicians and journalists.

Hacktivism

Belarusian hacktivists continue hack-and-leak operations against the Minsk regime.

Threats in the World

China: China is establishing full control over all domestic knowledge of software vulnerabilities. As always, China is active on social media, working to amplify pro-Chinese messages.

Russia: Political opposition and anti-corruption entities in Russia fall victim to DDoS attacks and data leaks. Stricter internet controls and censorship established before the September parliamentary election remain in place after the election. Proposed legislation prohibits foreign companies from processing biometric data of Russian citizens.

Iran: Iranian governmental websites were taken offline after a "cyber disruption".

North Korea: A North Korean cyber threat actor compromised a major South Korean major producer of combat ships & submarines.

Figure 30 Example of metrics and narrative, credit CERT-EU

Annex 3: Sample report

Development of the threat landscape

Who?	Group / Malware?	Why?	Trend
Adversary 1	APT-X	Adversary known to steal intellectual property in high tech industry.	↗
Adversary 2	APT-Y	Adversary known to steal intellectual property in our sector.	→
Targeted Cyber-Crime	FIN11 (TAS05)	Public and corporate IT-infrastructure is a growing market for ransomware	↗

Notable incidents and threat developments

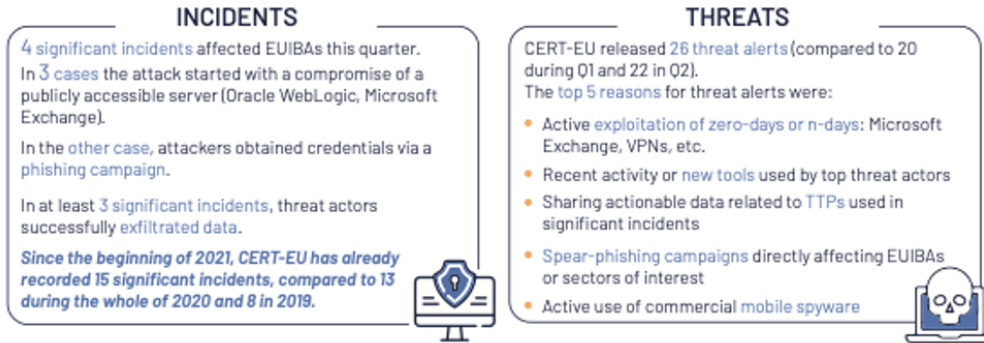



Figure 31 Credit CERT-EU

Coverage of key controls



Impact of additional measures on mitigation of the cyber risk

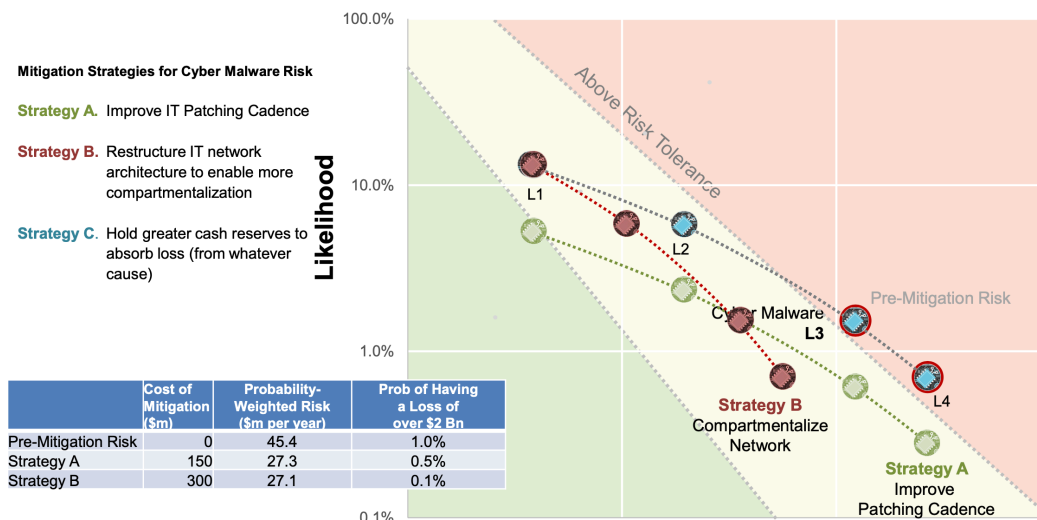


Figure 32 Credit Center for Risk Studies, University of Cambridge