Reporting Cyber Risk to Boards

Ten key insights for informed cyber oversight

Authors
Freddy Dezeure
Peter Debasse
João Pedro Gonçalves
Tristan Guiheux
Éireann Leverett
Patrick Mana
Lokke Moerel
Bartosz Sygula

Reviewers
Greg Bell
Paolo Borghesi
Philippe Coffyn
Chris Deverell
Tom Gilis
Kevin Holvoet
Angelos Keromytis
Ed Millington
Dimitri Rombaut
Sam Singer

Date: 30 August 2023
Version: Final

Table of Contents

INTRODUCTION 3

TEN KEY INSIGHTS 4

1. **Evidence rather than compliance** 4
2. **Reporting KCIs rather than everything** 4
3. **Threat-informed rather than stale** 5
4. **Priorities rather than averages** 5
5. **Reporting gaps rather than "all green"** 6
6. **Embedded rather than disconnected** 6
7. **Transparency of deviations rather than acceptance** 6
8. **Risk appetite rather than zero risk** 6
9. **Telling the story – risk connection to services** 7
10. **Unify cyber regulations – apply selective 'gold-plating'** 8

THE CISO'S REPORTING LINE(S) 8

PRODUCT, PORTFOLIO, SUPPLY CHAIN CYBER RISK 9

COMPARING WITH PEERS 9

Introduction

In March 2022, we published the white paper Reporting Cyber Risk to Boards, providing guidance for Chief Information Security Officers (CISOs) to design and implement quantitative cybersecurity metrics to report cyber risk at Board level and to provide reasonable assurance that the cyber risk is within accepted risk appetite. The white paper received a lot of attention and credit in the community and has been disseminated broadly. The white paper was also released in a condensed version for Board Members.

Since the publication of the white paper, additional regulatory requirements in the EU (NIS2[1], DORA[2]) and the US (SEC[3], NYDFS[4]) have increased the responsibility and accountability for Board members for exercising careful and informed oversight of cyber risk in their organizations. Cyber risk also plays an increasing role in ESG reporting. Some of these regulatory requirements make explicit reference to cyber metrics (DORA, article 6). There is currently no official guidance yet on what would constitute due oversight by Boards, let alone which strategic metrics could lead to *informed* oversight.

Feedback from the community on the content of the 2022 white paper and additional insights have indicated a need for additional guidance to highlight the main lessons learnt and tighten their formulation. This paper aims to serve this purpose. It also provides necessary elements to fulfill the additional regulatory requirements on Board information and oversight.

This paper builds on the fundamental notion that sound cyber risk management should be *evidence-based* rather than relying on intentions or assumptions (often based on self-reporting). Strategic cyber metrics are an essential component of any successful effort to prioritize and implement cyber risk management.

Measuring cybersecurity risks in a *quantifiable* manner, using data from the infrastructure, is not yet widely adopted across the industry. Unsurprisingly, there are also no agreed yardsticks to compare across peers.

The current paper aims to share 10 key insights from organizations that have implemented strategic cyber metrics so that the community can build on these learnings and adopt them in their own environment. The insights are summarized in crisp and thought-provoking summaries to facilitate uptake.

This approach may come across as a generalization or simplification, but it will help to streamline and focus the attention of CISOs and Boards and ultimately achieve better results in cyber risk management and oversight.

The current paper is to be considered as a companion to the 2022 white paper and it is highly recommended to read them together.

[1] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022L2555 (Articles 20 and 21)
[2] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022R2554 (Articles 5 and 6)
[3] https://www.sec.gov/news/press-release/2023-139
[4] https://www.dfs.ny.gov/system/files/documents/2022/10/rp23a2_text_20221109_0.pdf

Ten key insights

1. Evidence rather than compliance

Many organizations have adopted a cybersecurity framework (NIST, ISO, CIS) or followed specific regulation and standards (PCI-DSS, Solvency/Basel II), combined with external audits and certification. This is often a requirement in their activity, for regulatory or business (insurance, customers) reasons. However, this approach should be considered as a baseline rather than a panacea. It provides standards and policies to adhere to, but it does not necessarily reflect the actual fitness for purpose for the specific business risks.

Are the agreed controls sufficient to lead to effective risk mitigation if functioning correctly? Are they deployed completely? Are they functioning as intended?

The most mature organizations use (continuous) evidence from data collected in their infrastructure to ascertain the effectiveness of their controls rather than rely on human assessment, self-reporting and questionnaires completed once a year. The required data collection and retention is, admittedly, an important challenge but these organizations consider that it is worth the effort. Professional judgement can remain a component in providing context in reporting to the Board, if evidenced by metrics derived from operational data sources and security tools.

2. Reporting KCIs rather than everything

Boards want oversight on what matters and not all controls in cyber security frameworks are equally important. Guidance from most frameworks is that a limited number of controls makes the largest impact on the risk mitigation.

Reporting Key Control Indicators (KCIs) and their evolution over time to the Board therefore makes more sense than aiming to report on all controls. That doesn't mean that the CISO loses sight of all the other controls, it sharply brings into view what makes the most impact in risk mitigation for a specific organization at a specific time.

Avoid confusing the concept of Key Performance Indicators (KPIs) with KCIs. A CISO should be interested in his/her team performance, but KPIs will not always be relevant to how well the cyber risk is mitigated. Board reporting requires indicators which are strategic of nature, and which are representative for the overall internal control environment and underpin their risk appetite. Detailed guidance on KCIs and their effectiveness and coverage can be found in the 2022 white paper "Reporting Cyber Risk to Boards – CISO Edition".

Here is a sample list of KCIs as a starting point:

KCI 1	Asset Inventory[5]	% assets in the inventory within policy
KCI 2	Privileged accounts	% privileged accounts managed within policy
KCI 3	Timely patching	% high risk patches within N hours # of known exploited vulnerabilities detected
KCI 4	Back-up	Maximum time to recover key assets (% of critical assets recoverable in N hours)

[5] An accurate and complete asset inventory is critical as it's the denominator for many of the KCIs

KCI 5	Endpoint protection	% endpoints configured in line with policy
KCI 6	Logs collection	% critical systems onboarded to log collection
KCI 7	Network security	% compliant key network security configurations
KCI 8	Third Party compliance	% compliant key third-party connections
KCI 9	Identity management	% coverage of systems using MFA
KCI 10	Major Incidents	% major cyber incidents with business impact
KCI 11	Risk Acceptance	# risk accepted policy deviations
KCI 12	Internet exposed assets security coverage	% of Internet exposed assets covered by security monitoring and regular security assessment
KCI 13	Crown jewel coverage	% of crown jewels covered by security monitoring, vulnerability scanning and regular security assessment
KCI 14	Origin of Security Incidents	% of security incidents related to failures from at least one Key Control Indicator

3. Threat-informed rather than stale

The threat landscape is evolving and our controls and KCIs should do so, too. Adversaries adapt their tactics and techniques to bypass our defenses. They are often better aware of our infrastructure and control gaps than we are. They monitor vendor vulnerability disclosure and react to them with shorter lead times than we do. Some have sufficient resources to buy sophisticated exploits.

To avoid negative impact on our business, we need to adapt our controls to the threat, taking our specific environment and assets into account. This requires understanding of the adversary tactics, techniques, and procedures, prioritization and realignment of controls, and continuous monitoring for indicators of behavior and compromise. Important developments deserve to be tracked and reported to the Board. And, of course, given due priority at the technical level.

4. Priorities rather than averages

To keep focus on what really matters also means that we must be careful with averages. By averaging out, deviations from critical controls may remain under the radar and outliers remain unreported. Therefore, we would recommend not to make an aggregate of the results of all the hundreds of controls that you have identified. It may be attractive from the engineering perspective to identify a percentage of coverage of the whole framework, but this averaging out tends to hide the main issues.

Similarly, averaging out within a specific control may hide important risks. If, for example, an organization aims to patch critical vulnerabilities within three days, medium risk vulnerabilities within a month and all the rest within three months averaging the patching performance could hide the most critical ones.

Report averages only when it makes sense for a specific KCI. More detailed guidance on KCIs and coverage can be found in the 2022 white paper Reporting Cyber Risk to Boards – CISO Edition.

5. Reporting gaps rather than “all green”

It’s totally fine to report the actual situation to the Board, including gaps to close. They need to hear this, if this is the reality. It will also help the organization to comply with regulatory oversight and underpin prioritization of investments.

When reporting gaps to the Board it is necessary to explain what risk they entail, and which measures are proposed to resolve them in a projected time frame.

6. Embedded rather than disconnected

The impact of reporting cyber risk to the Board will be increased by providing access to the status of the controls to those who manage them (operators, managers). We call this the “democratization of the metrics”.

Reporting cyber risk to the Board inherently drives the organization. What is reported as being important will inevitably (fortunately) be perceived as important by the Board and within the organization. The KCIs reported should therefore make sense from the risk management perspective and reveal the true status of the risk.

The underlying data for the KCIs should be collected from the systems implementing the controls. Implementing connected metrics dashboards at all levels of the organization, with the granularity required to provide insight to the managers of the controls, creates transparency, increases ownership, and allows to finetune the system.

7. Transparency of deviations rather than acceptance

You may want to give visibility to deviations from key controls by reporting them to the Board. These deviations could stem from risk acceptance or policy violations (deliberate or inadvertent).

Most organizations have a process in place that allows departments to deviate from security policies by “accepting the risk”. Rather than keeping these deviations under the radar, it would be wise to report them. Giving visibility to deviations could help the organization to align to the key controls that are designed to mitigate the risk and stay within the risk appetite.

Monitoring deviations is especially useful in understanding maturity of the organization with risk management processes and culture. More mature ones tend to treat “risk acceptance” as the last of the available options, not the first.

The very process of documenting these deviations also allows us to identify any thresholds that are unrealistic, for example patching all vulnerabilities with a budget of 1% of ARR. By discussing the deviations, the whole organization can move towards risk acceptance thresholds that are more practical and realistic.

8. Risk appetite rather than zero risk

We can’t repeat this enough, an organization needs to determine at the Board level what an acceptable level of cyber risk is. Zero risk is an impossible and probably even undesirable target. It is essentially risk avoidance rather than ‘efficient treatment of risk’. In many organizations this risk appetite has already been established within the general business risk processes.

If this is not yet the case for cyber, the CISO should prompt the Board to determine the cyber risk appetite:

- How much are we willing to lose in the event the cyber risk materializes? Think about days of downtime, IPR theft, PII loss, reputational damage…
- To what extent do we want the risk to be mitigated? Zero risk is an impossible target. The cyber risk appetite is expected to be modulating between high and medium, given the evolution of the threat landscape and available technology.
- Which resources/budgets are we willing to make available for mitigation?
- Do we want to insure or self-insure the residual cyber risk?

A quantitative approach to cyber risk appetite is currently difficult to accomplish and is the exception rather than the norm. The Risk Appetite Statement (RAS) is usually based on a qualitative approach combining underlying qualitative and quantitative elements. RAS level is set by the Board in terms of Low, Medium, High. Often it is set by comparing different risk domains and prioritizing them. It is rather a calibration of the different domains. The actual underpinning of the RAS is a real challenge. It requires a cascade of indicators starting from the technical/operational level into the managerial and strategic level.

Though difficult, discussing cyber risk in terms of business impact in numbers is useful and the aim is not perfection. Be wrong about these numbers first, and let the executives work toward answering these questions in a repeatable manner. They should become aware of cyber as a business risk and deal with it in a similar way.

9. Telling the story – risk connection to services

A CISO needs to tell the cyber story in a business context for the message to be received. This requires him/her to understand the status of the controls, and their impact on the risk profile that is driven through business services.

An important objective for organizations that aim for true maturity of their risk and control environment is the capability to understand:

- how their business services (e.g. provide a loan or perform trading) impact the cyber inherent risk profile (e.g. the need to securely store restricted client data)
- and vice versa: how the inherent cyber risk can impact these services (e.g., lack of secure storage for restricted client data may result in inadvertent data disclosures or making data access much easier for malicious actors).

The organization must make sure that they understand which IT assets and processes support their Business Services (e.g. which systems are necessary to provide a loan). This allows to profile and measure the cyber inherent risk.

Next step is to ensure that cyber controls are applied and implemented to the IT assets and processes via automated methods (e.g., “controls as code”), so the effectiveness can be measured clearly via key control indicators (e.g. as mentioned earlier KCI5 % endpoints configured in line with policy).

This approach allows the CISO to provide the Board with a clear story about how current residual risk status and known control deficiencies can impact the business and give context to successfully manage risk appetite and make relevant investment decisions.

10. Unify cyber regulations – apply selective 'gold-plating'

Most contemporary and emerging cyber regulations globally have overlapping requirements. A CISO should implement all regulations pertinent to his/her organization in the different geographical locations and sectors by using mappings to unify their implementation.

This approach allows for the application of the same logic and rationale when addressing similar risks and controls, irrespective of the regulation in focus, with significant reductions in overhead for the CISO and the technical teams.

However, exceptions invariably exist. Some regulations may stipulate specific controls that are unique to a given entity and potentially challenging to maintain. In such instances, one can opt for a 'gold-plating' strategy, isolating these controls, exclusively for that entity. This selective application minimizes superfluous work for other entities within the organization and allows for maintaining a globally comprehensive cybersecurity strategy and report.

The CISO's reporting line(s)

A CISO is tasked with setting and maintaining the organizational vision, strategy, and program to safeguard information and technology assets. The CISO should be able to act autonomously and independently (e.g. DORA art. 6.4 [6]). Traditionally, the CISO reports to a C-suite executive, such as the Chief Information Officer (CIO), Chief Operating Officer (COO), Chief Finance Officer (CFO) or even the CEO. While this structure is widely adopted, CISOs may find themselves in a conflicting position where the relevant C-suite executive is also responsible for other functions which involve decisions on trade-offs between compliance with security standards and operational efficiency etc.

To avoid that the CISO acts in isolation, we need to maintain a balance between the interests of the internal stakeholders and mitigating the cyber risk. An organization may want to establish an Information Security Steering Committee (SteerCo) with the mandate to take operational decisions, monitor security risks and key controls, agree on metrics, sanction budgets, validate the security strategy and monitor its effective implementation.

The effectiveness of the SteerCo hinges on the participation of relevant C-suite members, such as the Chief Risk Officer (CRO), Chief Operating Officer (COO), Chief Compliance Officer (CCO), Chief Information Officer (CIO), Chief Financial Officer (CFO), Legal Counsel and of course, the CISO. A less frequent but empowered SteerCo is preferred over frequent SteerCo meetings with limited decision power.

The reporting of cyber risk to the Board would be in the remit of the CISO, ideally in agreement, or at least in full transparency, with the SteerCo. The CISO should

[6] https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32022R2554 (Article 6)

have an independent reporting line to the Board or one of its subcommittees, like the Audit Committee. The frequency of reporting of cyber risk to the Board should be commensurate to the materiality of the risk to the organization, but a quarterly report would be a good practice, if combined with an escalation process in case of need.

This model combines effective decision-making power as well as robust and effective governance.

Product, portfolio, supply chain cyber risk

The principles described in our white papers for reporting enterprise cyber risk to Boards can easily be transposed and extended into *product cyber risk* (how well are your products protected?), *portfolio cyber risk* (which are the key controls you would like to impose on your portfolio companies and how do you want to measure adherence to these key controls?) as well as *supply chain risk* (key assets, dependency, key controls and how to measure and report adherence). The KCIs may differ in these areas, still using similar principles.

Comparing with peers

We have found a substantial level of alignment on the principles outlined in this white paper within a cross-sector community of forty organizations that gathered in a CISO working group on a quarterly basis over a period of two years.

We hope that sharing these practices within the broader community will open the way to compare notes (and results) with peers within the sector and even to use these principles in the interaction with regulators.