Practical Cyberrisk Management

Cybersecurity and information security (cyber) professionals, staff, managers and board members are faced with the real issue of how to deal with cyberrisk, and this is a high priority for board members. There are many theoretical cyberframeworks and standards available, and many vendors sell "silver-bullet" cybertools that make practical cyberrisk management seem like a complex and daunting task. Applying basic and interrelated (framework and tool-agnostic) principles on what to protect, how to protect it and how to report on it to the board can make cyberrisk more manageable and contribute to a reduction in risk for the enterprise as a whole.

> "CYBERPROFESSIONALS PROTECT THE ORGANIZATION FROM A CYBEREVENT ON THE ONE HAND AND, ON THE OTHER HAND, GIVE ASSURANCE TO THE BOARD THAT THE ORGANIZATION IS PROTECTED."

Cyberrisk is about the impact to the business as a risk type alongside other risk, e.g., credit risk where clients do not pay their debts, product risk where a new product line fails or is of inferior quality, or strategic risk where the organization fails to achieve its strategic goals. It is part of operational risk and contributes to reputational risk. The impact of cyberrisk can be significant to the point of being terminal to the organization, and the risk must be managed.

Cyberrisk must be viewed from a business point of view as it is easy to fall into the trap of thinking that cyber is about the settings of a UNIX server or the rules on a firewall. While these things are important, it is a very narrow view and too far down the value chain to be used as a starting point. Cyberrisk is ranked among the top risk factors about which boards are concerned.[1, 2] A cyberevent can be a trigger that causes other business risk types to materialize, e.g., a liquidity event due to a run on a bank after a cyberevent. The impact on the business can be material and can be so severe that a well-managed and successful organization that is profitable and running smoothly can be severely damaged or terminated by a single cyberevent. Business, the board in particular, wants to know what the potential impact is of all the risk the organization faces, and that includes cyberrisk. The board wants to know if enough is being done at a fast enough pace to mitigate cyberrisk sufficiently. This point of view, therefore, must be the starting point of how to view cyberrisk. Cyberprofessionals protect the organization from a cyberevent on the one hand and, on the other hand, give assurance to the board that the organization is protected. Cyberrisk management aims to understand what needs to be protected, how to protect it, how to report on it to the board and how to indicate if the

Jaco Cloete, CISA, CRISC, CISM, CA, CCISO, CISSP
Has 22 years of experience in cyberrisk management and auditing in the banking sector. He can be reached at *https://www.linkedin.com/in/jacocloete*.

pace of achieving cyberresilience is fast enough. It is critically important that the board understands what it is being told.

Board Reporting

Reporting to the board can be an intimidating task. It is easy to fall into the trap of producing board reports containing technical jargon and statistics about which board members do not care. A board report must be kept short (usually limited to one or two pages), to the point and be visual.

Normal board reporting must address the critical questions that board members ask. The following aspects are typically enough to include in a normal cyberrisk board report:

- Cyberrisk posture in the form of a graph or a dial
- Progress with improving cyberrisk maturity
- Progress on cyberprojects
- Top five to 10 key cyberissues the organization is concerned about and what is being done to address the risk
- If actual vs. budgeted expenditure is on track
- Overall conclusion on whether progress to address cyberrisk and improving cyberresilience is fast enough

Sometimes, the board requests a special state of cybersecurity type of report, which can then be a lengthier report, but it must still be jargon and acronym free. The report must tell a story about the cyberjourney through headings, headlines, first sentences, powerful visuals and call-outs. Any detailed technical feedback must remain in the annexures.

Cyberrisk Posture

The crux of cyberrisk management is to harness the data in the organization and utilize it to tell a story about the cyberrisk posture. This is achieved by harvesting data from relevant data sources and building data models that feed cybermetrics which, in turn, are aggregated to an overall cyberrisk posture for the organization and for the individual business areas.

> “THE CRUX OF CYBERRISK MANAGEMENT IS TO HARNESS THE DATA IN THE ORGANIZATION AND UTILIZE IT TO TELL A STORY ABOUT THE CYBERRISK POSTURE.”

The better the data that are harvested and used, the better the quality of information available for decision-making. Data can be seen as interconnected, building a multidimensional model of the cyberuniverse. The following list indicates examples of source data that can be collected to form attributes of the various models that will be created:

- Log on data to identify a workstation ID and the employee logging on the workstation (workstation model)
- Encryption status of the workstation's hard drive (workstation model)
- User details, e.g., name, surname, employee number, business division and area (user model)
- Software and versions installed on a workstation (workstation model)
- User cyberawareness score (user model)
- User application access risk score (user model)
- Users with elevated privileges (user model)
- Users with universal serial bus (USB) access (user model if managed through user group policy object)

The following principles can be applied to create the models:

- Build a model for each type of item that is of importance. In this example, two models can be built: one for a workstation and one for a user.
- Extract the data from the data sources that can provide the previously mentioned data via an

application programming interface (API), service account or extract, transformation, load (ETL) tool, and land the data in the appropriate repository.

- Join the various tables on common keys and build a workstation model with its attributes and a user model with its attributes. Any data source items that cannot join to the other table become reconciling items. All reconciling items must have a reason, and all reconciling items without a valid reason become exceptions. Exceptions must be sent back to the data source owner to fix the data quality at source, e.g., to load a missing agent or to enroll an asset into Active Directory, or it could even be an indicator of compromise (IOC) that must be investigated and managed to resolution and redeployment to a precompromised state.

  An example of an IOC could be a workstation where the antivirus engine or a software management agent was removed by the attacker. The workstation will exist in the workstation model from other data sources, but there will not be an attribute for the software agent data source.

Once individual models have been developed, different models can be joined, e.g., joining the user model to the workstation model using the employee number as a common key. Linking the business division and area where the user works via the user model enables the attribution of noncompliance to business area and assists with various supplemental issues, e.g., fixing the asset register because the actual user of a workstation is not necessarily working in the same area in the business where the workstation was initially captured.

The main use of the models, however, is to measure noncompliance and attribution thereof. The models provide drill-down capability to assist business areas with granular reports for remediation purposes. Once the models are built, it is important to schedule the refresh interval. The closer the refresh interval is to real time, the greater the value, but the more expensive it becomes to maintain and operate.

The models and exceptions can be linked to the security information and event management (SIEM) system as an added benefit. Once the models have been built, create a set of metrics and risk indicators that are fed by the data from the models. The metrics can be coded into a data visualization tool that references the cyberdata models. An example of a metric is the percentage of unencrypted workstation hard drives. The data for this metric can be extracted easily from the workstation model by filtering all the workstations of which the encryption indicator is negative and dividing that by the total number of workstations in the workstation model.

Thresholds must be set for each metric. For granular reporting, each business area can set its own thresholds based on its business requirements. For one business area where workstations are used in a physically secure area with tight controls, it might be acceptable if workstation hard drives are mostly unencrypted, and this business area might choose an upper threshold of 20 percent and lower threshold of 10 percent. Results above the upper thresholds are red, results below the lower threshold are green and in-between results are amber. For another business area where users travel with notebooks, it might be unacceptable to have any unencrypted hard drives, and both the upper and lower thresholds are 0 percent. For a consolidated organizationwide view, the thresholds can be weighted based on contribution per risk driver, population per business area and by importance of the control being measured to the reduction of risk.

The consolidated view becomes the overall cyberrisk posture ideally depicted as a dial, as per **figure 1**. Individual business areas can be depicted by a dial for metrics specific to that business area.

Figure 2 indicates an example of how models can be created in principle by joining source data tables. From this example, the workstation model consists of the logon data and encryption data, which are joined with join 1. Human resources (HR) data form the base for the user model, and the two models are joined with join 2, indicating that workstation WS10019293 is not encrypted and belongs to the retail sales area as John Doe (EMP12345) logs on to it mostly.

> “IT IS VERY IMPORTANT TO IDENTIFY CROWN JEWEL ASSETS BECAUSE NOT ONLY WILL THIS IDENTIFICATION PROVIDE INSIGHT TO THE CYBERPROFESSIONAL ABOUT HOW THE ORGANIZATION OPERATES, BUT IT WILL ALSO GUIDE THE ORGANIZATION ON WHERE TO FOCUS CYBEREFFORTS AND RESOURCES.”

A practical way to build data models is to apply dimensional modelling in a data warehouse by utilizing a star schema, dimension tables and fact tables.[3] Data from source systems, e.g., encryption status, become attributes of dimension tables, e.g., workstation model and user model. Aspects that need to be measured become fact tables created by joining dimension tables to the fact tables.

Crown Jewels

The important question is which metrics to choose and how to differentiate between the importance of different metrics. This is achieved through crown jewel identification[4] and threat modeling. Crown jewels are those critical assets that, if exposed, could cause significant and potentially terminal harm to an organization.

It is very important to identify crown jewel assets because not only will this identification provide insight to the cyberprofessional about how the organization operates, but it will also guide the organization on where to focus cyberefforts and resources. This daunting task can be achieved by following these steps:

- Have planning sessions with the business information security officer (BISO) and risk officer to understand the enterprise better and understand how the organization is structured.
- Agree on the segregation of the enterprise in logical groups based on function or location.
- Set up workshops of two to three hours with each group. In the workshops, discuss the business processes and keep asking questions such as what could cause big harm to that business area and what systems and processes, if impacted, would cause the enterprise to not function anymore. Several types of impacts can be considered (e.g., How do money and data flow through the business? Are payment instructions generated from the area? Are there

Figure 1—Example of an Overall Cyberrisk Posture Dial

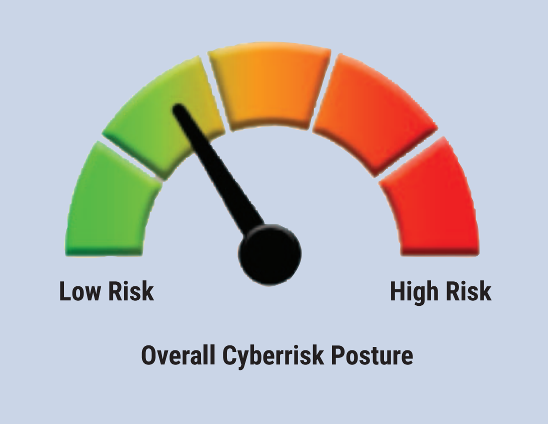

Figure 2—Joining Data Sources Example

Source 1 (Log On Data)			Source 2 (Encryption Data)			Source 3 (HR Data)					
Rec No	Log On ID	Workstation Name	Rec No	Workstation Name	Encryption HDD	Rec No	Emp ID	Name	Surname	Business Division	Business Area
1	EMP12345	WS10019293	1	WS10019293	No	1	EMP12345	John	Doe	Retail	Sales
2	EMP10473	WS10018939	2	WS10018939	Yes	2	EMP10473	Susan	Jones	HR	Training
3	EMP12123	WS10039293	3	WS10039293	Yes	3	EMP12123	Rafik	Said	Finance	Budgeting
4	EMP13245	WS10022343	4	WS10022343	Yes	4	EMP13245	Jacques	Le Roux	IT	Server ops
5	EMP11123	WS10023423	5	WS10023423	Yes	5	EMP11123	Herbert	Schmidt	Risk	Audit

Join 1—Workstation Model

Join 2—Workstation and User Model

sensitive files managed in the business area? Are there regulatory requirements in the area? And, are there systems with sensitive functions [e.g., systems affecting payments])?

- Note all the answers, but afterward distill iteratively until there are one or two key systems/processes listed for the business area.
- Repeat this process until all the key systems and processes in all the business areas have been covered.
- Afterward, distill the combined lists iteratively until there are no more than 10 critical systems/processes left. In a bank, for example, it could be the SWIFT system and supporting processes or the payment switch, and in a manufacturing organization, it could be the system managing the plant operations. And in a retail operation, it could be the client repository.
- Systems can be a crown jewel based on function. If in use within the organization, Active Directory and domain controls systems could be crown jewels, as well.
- Validate the final list of important systems and the ones marked as crown jewels with the business risk officers, BISOs and senior management. The cyberprofessional is then armed with the list of crown jewels to inform the other cyberrisk management processes.

Threat Modeling

A cyberprofessional must determine what are the worst things that can happen to the organization. This can be achieved by analyzing the list of crown jewels and determining how an attacker would pull off a successful attack against each crown jewel.

Figure 3—Crown Jewels

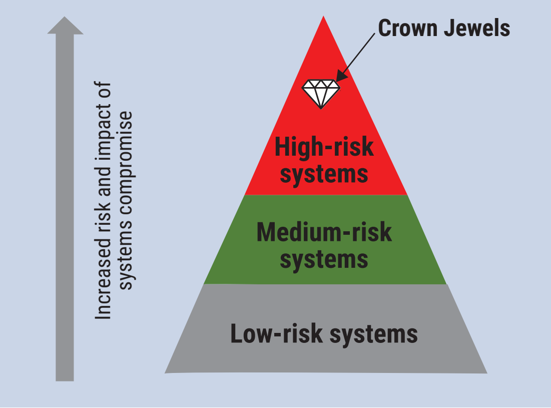

A practical way to do this is through attack path mappings. Take each crown jewel and analyze the ways that the crown jewel can be compromised. This can be executed by in-house skills or through contracting with a reputable niche cybersecurity company. It is important to consider threat modeling from a holistic point of view, which requires a mind shift by looking at cyber holistically. Some organizations take a very narrow view of cyber. The following example explains the concept.

> “IT IS IMPORTANT TO CONSIDER THREAT MODELING FROM A HOLISTIC POINT OF VIEW, WHICH REQUIRES A MIND SHIFT BY LOOKING AT CYBER HOLISTICALLY.”

Suppose a cyberpractitioner contacts one of the organization’s third-party suppliers with whom data are exchanged to obtain assurance over the security in the third party’s environment. The third party is only willing to discuss security relating to the particular server and network segment used for service to the organization. This is not the correct approach. Further, assume the third party accesses the network segment via a jump box on the third party’s main network, and that jump box is managed by Active Directory on the same network as a workstation in a remote location of the third party’s network. This means that if the remote workstation is compromised at the third party and the attacker manages to move laterally on the network and do privilege escalation up to the point where the domain is compromised or where a jump box user’s credentials are obtained, the jump box will be compromised and all segments accessed by the jump box will be compromised, and it will be game over. In this case, the whole environment at the third party is in play and not only a particular segment, hence the importance of taking a holistic view of cyber.

A very handy tool to use in attack path mapping is the cyber kill chain developed by Lockheed Martin.[5] With this model, an attack can be dissected and controls considered through each stage of an attack. What becomes clear after a while is that many attack paths occur via the domain and, therefore, the domain controller becomes a crown jewel due to its function and will most probably be a crown jewel in most organizations. Examples of attack paths that are universal in most organizations are compromising a workstation through infection by malware for the attacker to gain access to the network to do further privilege escalation to crown jewel X and compromising a user via a phishing attack to obtain his or her credentials to gain access to the network to do further privilege escalation to compromise crown jewel Y. Focusing cyberefforts on workstations and users is a very good start.

Using the Crown Jewel List and Threat Modeling Attack Path Maps

The crown jewel list and initial attack paths become powerful tools for the cyberprofessional that can be used to:

- Enhance the cyber management information system (MIS) by informing which metrics should be measured and tracked along the most probable attack paths.
- Enhance reporting by focusing on reporting of cyberrisk related to the crown jewels.
- Guide red team and purple team testing.
- Motivate funding and budgets.
- Inform business impact analysis for business continuity management.

The cyberpractitioner, together with the business representatives (e.g., BISO and risk officer) must unpack the crown jewel list into its components. The idea is to determine all the individual components that form part of a crown jewel so that a data source can be created that can be linked to the existing cyberdata models. The components might consist of the workstations that are relevant to the crown jewel, the users who use and log on to the crown jewel, the servers relevant to the crown jewel, and any other components that are relevant to the attack path for which there are existing data models. These components can be manually populated into a spreadsheet by all the business areas and a table of the spreadsheet joined via the common key of the component, e.g., employee number, server name or workstation name. Linking the crown jewel components to the data models enables the cyberprofessional to produce granular, crown-jewel-related reporting to get a cyberposture only on crown-jewel-affected items. **Figure 4** indicates an example of a data table consisting of the components of the customer repository crown jewel.

Expanding on the first example depicted in **figure 2**, if the data table in **figure 4** are joined to the existing data models, it can be seen that workstation WS10019293 is a component of the customer repository crown jewel and receives a crown jewel indicator in the workstation model. From **figure 2**, it can be seen that it is used by user John Doe from retail sales and will most probably contain extracts of client data. From the workstation model, it is evident that the hard drive is not encrypted. Tying this information back to risk appetite, the high-level board statement could be that there is no appetite for loss of concentrations of client data. Users and workstations are drivers for that statement, and measuring the metric of unencrypted workstations,

Figure 4—Crown Jewel Components

Rec No	Crown Jewel Name	Component Type	Component Name
1	Customer repository	Workstation	WS10019293
2	Customer repository	Server	CIRIUS_SQLPRD
3	Customer repository	User	EMP12345
4	Customer repository	Database	CLIENT_SQLPRD
5	Customer repository	Workstation	WS10022343
6	Customer repository	User	EMP13245

and due to the crown jewel indicator, the organization is operating outside of its risk appetite.

SIEM rules can be programmed and prioritized, focusing on the crown jewel list and crown-jewel-affected items. Special controls can be implemented on crown-jewel-affected items. Special controls are those controls that do not necessarily need to be implemented on all items on a network because of cost or effort, but are warranted to implement on crown-jewel-affected items. There should be fewer budget restrictions on special controls. In the previous example, a special control would be in-depth classroom training for John Doe and Jacques Le Roux, who are working on the client repository crown jewel, to make sure that they are equipped to look out for tailored attacks that could be used to pull off a breach of that crown jewel. These two employees' workstations would also be contenders for special expensive monitoring controls.

Analytics on attack paths produce patterns whereby certain components emerge as pervasive across attack paths and pervasive types of controls can be applied, e.g., if all users are configured to receive email, a comprehensive security awareness program against phishing attacks across all users can be implemented. Also, malware and lateral movement on workstations is relevant to all workstations and, therefore, endpoint detection and response (EDR) software can be rolled out across the base.

Red Team Testing

Red team testing provides insights to how well the organization's technical team can defend against the latest real-life attacks. It provides deep insights into blind spots across the kill chain.

The Bank of England implemented a program (CBEST framework) where certain financial organizations are subjected to red team testing by Council of Registered Ethical Security Testers (CREST)-certified security enterprises.[6] Recently, the European Central Bank also published the European framework for Threat Intelligence-based Ethical Red Teaming (TIBER-EU).[7] This same principle can be adopted and applied to any organization. A practical way to do this is to subject the organization to an attack by the attacking team, which is called the red team. The red team attack is open scope, and all attack paths are in play. Attacks can originate from the Internet to test the perimeter controls, but can also start on the internal network with an "assumed perimeter breached" approach. The organization must be defended by its security team, which is called the blue team. The attacking team always has the advantage, and given enough time, most organizations will eventually succumb to a persistent attack by a good red team. The blue team must deploy preventive and detective controls to try to prevent the red team from entering the network and moving along the kill chain or detect the red team and kick them out of the network.

> "ANALYTICS ON ATTACK PATHS PRODUCE PATTERNS WHEREBY CERTAIN COMPONENTS EMERGE AS PERVASIVE ACROSS ATTACK PATHS AND PERVASIVE TYPES OF CONTROLS CAN BE APPLIED."

This cat-and-mouse game has the advantage of highlighting weaknesses where there are blind spots. There will always be weaknesses because no organization is 100 percent secure. The testing also keeps the blue team battle fit and provides valuable experience to see real-life attack behavior on the monitoring systems. It is just a matter of time until regulators across the world pick up on red team testing and make it a requirement to have these tests done. It is in the interest of all organizations to implement red team testing whether required by regulation or not. The outcome of red team testing can be used to inform attack paths and can become part of the board reporting on top cyber-related items to address. Red team testing differs from traditional penetration testing whereby a penetration test is very narrowly focused on a specific target within specific boundaries; a red team test has a much wider scope where the red team tries various routes across the estate until the easiest path is found to achieve the assignment objectives. It makes sense to use the crown jewel list when creating the scope for a red team test.

Purple Team Testing

A valuable derivative of red team testing is purple team testing, where the red team and blue team work together as a purple team.

Purple team testing typically takes the form of the red team executing an attack while sitting in the same room as the blue team and the blue team checking if the attack was detected. If the blue team does not detect the attack, the red team shows the blue team what was done, and the blue team tweaks the SIEM rules until the attack is detected. An example of purple team testing is to produce verified attack paths for immediate remedial intervention. This is very expensive, but warranted if done on a crown jewel. A crown jewel is selected and attack paths on the crown jewel are analyzed in a white-box approach (all information provided by the blue team), where the red team applies the latest attack techniques to determine as many attack paths as possible to crown jewels through actual exploitation in real life and the blue team applies technical defensive controls as the exercise continues.

Response

At one point or another, even with the best and most diligent controls in place, an organization will be compromised to some extent. To be resilient, it is important that organizations prepare for when this happens by simulating cyberevents and testing the response to the event. This will provide valuable insight to the board members who participate and is a physical way of reporting on cyberrisk management through first-hand experience.

From time to time, organizations must simulate cyberattacks either through dedicated cybersimulations or as part of another type of simulation but with a cybertrigger. The testing includes responses to the media and notifications to key stakeholders, e.g., shareholders, regulators and the public. It is imperative for an organization to have a social media team that monitors and responds to social media comments regarding the organization, and this team must also be part of cybersimulations. It is also important to have top management involved in the simulations, including the executive level. Having the executives and board members involved in simulations has the advantage of the board getting comfortable with cybercontrols in place due to interaction with cyberteam members and the teams involved in the exercise to get to know the executives better.

> “AUDITORS MUST HAVE A MIND SHIFT TOWARD CROWN JEWEL AND ATTACK PATH THINKING WHEN PLANNING AUDITS.”

Auditing

While there is value in doing traditional audits focusing on, for example, databases and operating systems, there is much more value in auditing the controls along the kill chain utilizing the attack paths to a specific crown jewel.

There is a risk that auditors might fall into the groove of traditional auditing that focuses on the same old database, operating system and network device audits. Auditors must have a mind shift toward crown jewel and attack path thinking when planning audits. The crown jewel list and attack paths must be shared with audit so that it is very clear where the high-risk systems and key controls are. It is much more valuable to do an end-to-end test of a SWIFT attack than, for example, an operating system audit, which might be one small component in the end-to-end process. Audits should be focused on providing assurance over the controls that would prevent the worst-case events from happening, the controls that would detect the events and the response controls that would reduce the impact of the event. The crown jewel list and attack path mappings become valuable input for the auditor to understand how critical attacks on the organization could occur, and audit efforts can then be focused on those key controls.

Advanced Cyberanalytics

If multidimensional models are in place with all their attributes, it is possible to take things to the next level by adding network flow data to the model and joining by IP address. Augmented by the crown jewel list and crown-jewel-affected items, this can provide information on network behavior and patterns of network traffic that will start to appear and could reveal additional attack paths not previously known. Over time, normal behavior can be learned through machine learning, and activity outside of the norm can feed the SIEM and be investigated by the security operations center/network operations center (SOC/NOC).

The Assertions of Existence vs. Completeness

The assertions of existence and completeness are incredibly important in cyberrisk management. If someone asks for a list of items (e.g., workstations), a list is printed from a system. This is how many things in real life operate, from asset registers to user lists to network diagrams. These lists are mostly valid because, when people follow processes, items will be added to the list if the process works as designed. The list is used for many things from change control to auditing to planning to sampling population. The biggest risk for cyber lies predominantly on the completeness side. The items not on the list are the items that will most probably be used in a cyberattack and could include the workstation without an agent, the rogue device plugged into the network that is not on the domain, the user account added on the server, but not in Active Directory, etc. It is not what is on the list, but what is not on the list that is of extreme importance.

Why do severities occur after changes? The affected item was not on the list and, therefore, not tested. Why did the rogue device go undetected for so long? It was not on the list and an agent was not loaded on it. As a result, it is very important to refresh data models of the live environment on a frequency that is as close to real time as possible. All unexplained reconciling items after joining source data must be investigated as either items that are not on the list or items erroneously on the list.

> “FROM TIME TO TIME, ORGANIZATIONS MUST DO CYBERRISK ASSESSMENTS USING ONE OF THE AVAILABLE FRAMEWORKS TO PROVIDE A CHECKPOINT FOR THE STATE OF CYBERSECURITY.”

Cyberscenarios

Banks that must hold capital in line with the Basel Committee on Banking Supervision (BCBS) principles can benefit from cyberrisk management by utilizing the aspects discussed so far.

Cyberevents can have a material impact on a bank and contribute to operational risk. If the organization knows and understands its crown jewels and the impact that a compromise to the crown jewels can have on the organization, these impacts can inform the potential worst-case scenarios of a cyberevent. Detailed analysis of external events and losses experienced by other organizations relating to similar environments as the organization’s own crown jewels can inform potential losses that could occur if the event happens at the organization itself. Analysis of the risk drivers via threat modeling, the kill chain, and the outcome of red team testing and attack path mapping can inform the story line of a scenario and the likelihood of the event happening. Metrics from the cyber MIS can assist with understanding the contribution and attribution internally in the organization toward the capital for improved decision-making and where to focus efforts.

Risk Assessments and Maturity Assessments

From time to time, organizations must do cyberrisk assessments using one of the available frameworks to provide a checkpoint for the state of cybersecurity. Gaps identified can inform future efforts as part of the cyberstrategy. A program of cybermaturity improvement must be maintained, and each control implemented improves the maturity level, which forms part of board reporting.

Conclusion

Cyberrisk management can seem extremely complex, and cyberprofessionals may not know where to start. Applying basic principles can make the task more manageable and may include:

- Implementing a cyber MIS solution with data sources building multidimensional models of the live environment
- Producing meaningful, short and visually impactful reporting to the board on aspects that matter
- Producing a crown jewel list with its components and link to the cyber MIS
- Performing attack path mapping to understand how crown jewels can be attacked to focus efforts
- Testing resilience of the organization's cyberdefenses by performing red team testing
- Performing periodic risk assessments to serve as checkpoints and maintain an ongoing cybermaturity assessment program

Author's Note

The views expressed in the article are those of the author and do not necessarily represent the views of his employer.

Endnotes

[1] Deloitte, "CEO and Board Risk Management Survey: Illuminating a Path Forward on Strategic Risk," 2018, *https://www2.deloitte.com/us/en/pages/risk/articles/ceo-board-of-directors-risk-management-survey.html*

[2] North Carolina State University's Enterprise Risk Management Initiative and Protiviti, "Executive Perspectives on Top Risks 2019," USA, 2018, *https://www.protiviti.com/sites/default/files/united_states/insights/nc-state-protiviti-survey-top-risks-2019-executive-summary.pdf*

[3] Kimball, R.; *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, *3rd Edition,* John Wiley & Sons Inc., USA, 2013, p. 7-18

[4] Information Security Forum, "Protecting the Crown Jewels: How to Secure Mission-Critical Assets," *https://www.securityforum.org/tool/protecting-the-crown-jewels/*

[5] Lockheed Martin, "Cyber Kill Chain," *https://www.lockheedmartin.com/en-us/capabilities/cyber/cyber-kill-chain.html*

[6] Bank of England, "Financial Sector Continuity," *https://www.bankofengland.co.uk/financial-stability/financial-sector-continuity*

[7] European Central Bank (ECB), "ECB Publishes European Framework for Testing Financial Sector Resilience to Cyber Attacks," 2 May 2018, *www.ecb.europa.eu/press/pr/date/2018/html/ecb.pr180502.en.html*