

LIBRARY OF

CYBER RESILIENCE METRICS

Types

#	● Service	Port
1	● http	80
2	● domain	53
3	● ms-term-services	3389
4	● unknown	21320
5	● microsoft-ds	445
6	● snmp	161
7	● ms-sql-s	1433
8	● ssh	22

Contents

Preface

5

About the Framework

6

Metrics specification

8

Ability to avert social engineering
 Ability to engage threat intelligence
 Ability to address vulnerabilities
 Ability to handle cyber incidents
 Ability to resist malware
 Ability to resist system intrusions
 Ability to resist DDoS attacks
 Ability to protect credentials
 Ability to protect key assets
 Ability to measure and minimize damage

9

11

14

17

21

24

27

29

31

34

SRP Cyber Security

37

References

38

Colophon

© 2017 Participants in the
 Cyber Security Shared
 Research Program.

Editor

Richard Kerkdijk (TNO)

Contributors

Roger Lagarde (ABN AMRO),
 Tommy Koens and Sander
 Zeijlemaker (ING), Paul
 Samwel (Rabobank),
 Bert-Jan te Paske, Eldine
 Verweij, Richard Kerkdijk
 and Reinder Wolthuis (TNO).

Copyright

All rights reserved. No part
 of this document may be
 reproduced and/or publis-
 hed in any form by print,
 photoprint, microfilm or any
 other means without
 previous written permission.

Contact

wegwijzer@tno.nl

Preface

Present day financial services rely heavily on electronic channels and complex IT infrastructures. This setup makes it possible to carry out financial transactions with speed and efficiency, while offering business and residential customers a wealth of features. However, it also makes financial services susceptible to cyber attacks. Financial providers have therefore invested heavily in provisions that ensure an appropriate level of cyber resilience. But what is true cyber resilience and to which extent are current measures achieving it? And equally important: which capabilities or working areas require improvement and which effects can be expected from specific further investments (e.g. acquisition of a technical security solution or specific specialist training)?

Compelling questions such as these evoked a strong desire among financial institutions to measure and quantify the state of cyber resilience within their organizations. Since it was felt that traditional security metrics offer limited insight into the actual performance of cyber resilience provisions, an initiative was launched to jointly define a meaningful framework of cyber resilience metrics. This work took place in a collaborative cyber security research program featuring TNO, ABN AMRO, Rabobank, ING and Achmea (see back cover).

This booklet was compiled to share the framework of cyber resilience metrics with other organisations that seek quantitative appraisal of their cyber security capabilities. We hope our work will help you establish an effective metrics program and (more importantly) ensure that your organization maintains a solid cyber security posture.

About the framework

The framework encompasses 47 metrics that were consolidated into 10 core categories. Figure 1 depicts the top-level framework structure.




Figure 1: Core categories of cyber resilience metrics

As the figure shows, the overarching structure was based on the “cyber kill chain” as developed by Lockheed Martin [Lockheed]. The parties involved in this initiative felt that metrics for cyber resilience should reflect an organisation’s readiness for so called “targeted cyber attacks” (also referred to as “Advanced Persistent Threats” or APTs) and the cyber kill chain is a widely recognized model for attacks of this nature. Each category in the framework represents a specific ability that is considered instrumental for averting or handling targeted cyber attacks. The “ability to avert social engineering” (category 1), for instance, reflects the adequacy of employees’ responses when faced with social engineering techniques such as phishing. Such techniques are a key element in many targeted cyber attacks (e.g. for the purpose of reconnaissance). Similarly, metrics in the “ability to resist malware” category (item 5 in the figure) reflect the organisation’s ability to detect, contain and remediate malicious

software that is present or active in its technical infrastructure. Typical attack scenarios involve the use of malware in various stages of the cyber kill chain.

By focusing on cyber resilience abilities and the actual effects achieved through technical and organizational security measures, this initiative pursued a material step forward in measuring an organisation’s cyber security posture. Traditional security metrics tend to focus on the existence of security controls or the fulfilment of specific security requirements. A typical example is that many organizations assess the state of security awareness among employees by measuring the extent to which they have been subjected to (mandatory) security training. In itself, however, the fact that an employee has completed an e-learning module offers little assurance that he or she will exhibit appropriate behaviour when faced with an actual security threat. This initiative aspired to overcome this by defining metrics that reveal the actual status and performance (outcome) of cyber resilience measures, thus offering a more viable foundation for managing security operations or justifying investments. To ensure that the metrics in the framework are indeed meaningful, an analysis was made of 23 APT-type attack scenarios that actually occurred in the financial industry (in the Netherlands or elsewhere). Each scenario was characterised in terms of abilities needed to avert or handle it in various stages of attack. In turn, each ability was translated into one or several metrics reflecting its state or performance at a given moment in time.

For a more elaborate explanation of the framework and its underlying philosophy, we refer to [Measuring].

Metrics specification

This section presents the actual cyber resilience metrics encompassed in the framework. Each metric is specified according to the following format:

Mx. Title	
Definition	Formal definition of the metric
Purpose	Rationale of the metric in terms of the insight it provides and possibly applicable limitations.
Differentiation options	Possibilities to differentiate the metric by context or application area in order to refine the acquired insights.
Data sources	Processes or technical facilities that are likely to offer (some or all of) the data required to quantify this metric in actual practice.

The metrics are grouped according to the 10 core categories depicted in figure 1. As explained above, each category corresponds to a specific cyber resilience capability.

Ability to avert social engineering

M1. Resistance to illicit phone calls	
Definition	% employees that recognise and report social engineering when subjected to a phone call assessment.
Purpose	Indicates the degree to which employees are capable of a. recognising a social engineering scheme conducted via illicit phone calls b. exhibiting desired behaviour if such a social engineering scheme occurs A higher percentage equals better performance.
Differentiation options	Can be differentiated by employee position or function group, e.g. general population versus senior management versus system maintenance staff. Note: when doing so, it would make sense to also differentiate the degree of difficulty of social engineering simulations employed.
Data sources	Security helpdesk or similar notification point for (suspected) security incidents.

M2. Exposure to phishing schemes	
Definition	% employees that fall victim to a phishing scheme when subjected to an exposure test
Purpose	Indicates degree to which employees are susceptible to phishing schemes. A lower percentage equals better performance.
Differentiation options	Can be differentiated by employee position or function group, e.g. general population versus senior management versus system maintenance staff. Note: when doing so, it would make sense to also differentiate the content and degree of difficulty of phishing simulations employed.
Data sources	Outcome of exposure test (phishing simulation), e.g. collected by counting visits to simulated phishing website.

Ability to engage threat intelligence

M3. Resistance to phishing schemes

Definition	% employees that report phishing schemes when subjected to an exposure test.
Purpose	Indicates the degree to which employees are capable of exhibiting desired behaviour when subjected to phishing. A higher percentage equals better performance.
Differentiation options	Can be differentiated by employee position or function group, e.g. general population versus senior management versus system maintenance staff. Note: when doing so, it would make sense to also differentiate the content and degree of difficulty of phishing simulations employed.
Data sources	Security helpdesk or similar notification point for (suspected) security incidents

M4. Assessment of threat notifications

Definition	% of threat notifications that was analysed to assess relevance and (potential) impact
Purpose	Indicates the organisation's ability to consume (large volumes of) threat intelligence. A higher percentage equals better performance.
Differentiation options	– By actual intelligence source, e.g. FS-ISAC¹, FIRST², CIRCL³, vendor x, etc. – By nature of intelligence source, e.g. internal vs. public vs. private vs. community – By nature of threat notification, e.g. IoC⁴, vendor advisory, trend report, etc.
Data sources	Security helpdesk or similar notification point for (suspected) security incidents

¹ Financial Services ISAC, a community of financial providers that a.o. exchanges cyber threat intelligence amongst its membership, <https://www.circl.lu>

² Global Forum of Incident Response and Security Teams, <https://www.first.org>

³ Computer Incident Response Center Luxembourg (CIRCL), a national CERT that active cyber threat intelligence communities for various types of constituents

⁴ Indicator of Compromise

M5. Operational follow-up on threat notifications

Definition	% threat notifications that invoked tangible action (e.g. modify firewall rules or monitor on IoC) in operational security processes.
Purpose	Indicates the organisation's ability to translate threat intelligence into actual security enhancements. A higher percentage indicates that a larger portion of collected intelligence led to tangible follow up and thus equals better performance.
Differentiation options	– By actual intelligence source, e.g. FS-ISAC, FIRST, CIRCL, vendor x, etc. – By nature of intelligence source, e.g. internal vs. public vs. private vs. community – By nature of threat notification, e.g. IoC, vendor advisory, trend report, etc.
Data sources	– Threat intelligence platform (if in place) – Threat intelligence analyst (manual administration)

M6. Effectiveness of threat notifications

Definition	% operationalised threat notifications that ultimately resulted in tangible effect (e.g. an IoC sighting or blocking of traffic from illicit source)
Purpose	Indicates the organisation's ability to separate relevant from irrelevant threat intelligence. A higher percentage indicates that a larger portion of collected intelligence led to an actual (security relevant) effect equals better performance. Note: should be calculated as [# threat notifications that resulted in tangible effect] / [# threat notifications that invoked any follow up] (see previous metric).
Differentiation options	– By actual intelligence source, e.g. FS-ISAC, FIRST, CIRCL, vendor x, etc. – By nature of intelligence source, e.g. internal vs. public vs. private vs. community – By nature of threat notification, e.g. IoC, vendor advisory, trend report, etc.
Data sources	– Threat intelligence platform (if in place) – Threat intelligence analyst (manual administration) – Security monitoring tools or SOC⁵ analyst (sightings) – Firewall logs or firewall maintenance staff – System logs or system operators

⁵ Security Operations Center

M7. Timeliness of intelligence processing

Definition	Mean time (hours, days) elapsed between receiving threat notification and processing it (i.e. either discarding it or initiating follow up)
Purpose	Indicates the organisation's ability to promptly respond to threat notification. A low score equals better performance.
Differentiation options	– By actual intelligence source, e.g. FS-ISAC, FIRST, CIRCL, vendor x, etc. – By nature of intelligence source, e.g. internal vs. public vs. private vs. community – By nature of threat notification, e.g. IoC, vendor advisory, trend report, etc. – By time of day or week, e.g. office hours vs. nightly hours vs. weekends/ holidays
Data sources	– Threat intelligence platform (if in place) – Threat intelligence analyst (manual administration)

Ability to address vulnerabilities

M8. Coverage of vulnerability scanning	
Definition	% IT assets covered by automated vulnerability scans
Purpose	Reveals the reach of vulnerability scanning operations and thus the ability of the organisation to assess existence of common (known) vulnerabilities in its IT infrastructure. A higher percentage equals better performance. Note that this metric is contextual in nature, since it does not reveal actual effects achieved.
Differentiation options	– By network segment, e.g. office network vs. production server infrastructure – By asset type, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Vulnerability scanner – Asset database (CMDB)

6 Configuration Management DataBase

M9. Coverage of penetration testing	
Definition	% IT assets subjected to penetration testing
Purpose	Reveals the reach of penetration testing activity and thus the ability of the organisation to identify vulnerabilities in software configurations that a vulnerability scanner would typically not pick up on. A higher percentage equals better performance. Note that this metric is contextual in nature, since it does not reveal actual effects achieved. Moreover, this metric will be most meaningful if applied to a certain (limited) timeframe, e.g. the percentage of IT assets subjected to pentesting in the past 3 months.
Differentiation options	– By network segment, e.g. office network vs. production server infrastructure – By asset type, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Penetration testing procedures and reports – Asset database (CMDB)

M10. Exposure to common vulnerabilities	
Definition	% IT assets that were mitigated of significant vulnerabilities
Purpose	Indicates the extent to which common (known) vulnerabilities in the organisation's IT infrastructure were remediated, thus reducing exposure to common exploits and abuse scenarios. A higher percentage equals better performance (i.e. lower exposure).
Differentiation options	– By network segment, e.g. office network vs. production server infrastructure – By asset type, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Vulnerability management process – Vulnerability scanner

M11. Exposure to skilled intrusion attempts	
Definition	% penetration tests that resulted in high risk findings
Purpose	Indicates the extent to which a skilled intruder could invade or otherwise abuse the organisation's IT assets. A lower percentage equals better performance.
Differentiation options	– By network segment, e.g. office network vs. production server infrastructure – By asset type, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Penetration testing procedures and reports

Ability to handle cyber incidents

M13. Timeliness of incident detection	
Definition	The average time (hours, days) that elapsed between start and detection of a cyber incident.
Purpose	Indicates the ability of the organization to detect cyber incidents with sufficient speed. A lower value indicates that cyber incidents are detected faster whereas a higher value indicates that it will take considerable time to detect an incident, thereby providing more opportunity to an attacker to achieve his goals.
Differentiation options	– By incident type, e.g. DDoS⁷ vs. system intrusion vs. malware vs. phishing vs. ... – By service affected. A financial provider might for instance quantify this metric separately for internet banking, debitcard payments, etc. – By network segment affected, e.g. office network vs. production server infrastructure
Data sources	– Security monitoring systems (SIEM⁸, IDS⁹) – Incident management processes and tools – Incident evaluation reports – Forensic investigation reports

7 Distributed Denial of Service

8 Security Incident and Event Management solution

9 Intrusion Detection System

M12. Timeliness of vulnerability remediation	
Definition	Average lifetime (hours, days) of vulnerabilities identified through scanning or testing
Purpose	Indicates the degree of responsiveness to software vulnerabilities, i.e. the organisation's ability to resolve such vulnerabilities within an acceptable timeframe. A lower number equals better performance.
Differentiation options	– By network segment, e.g. office network vs. production server infrastructure – By asset type, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Vulnerability management process – Vulnerability scanner – Penetration testing procedures and reports

M14. Timeliness of incident mitigation

Definition	The average time (hours, days) that elapsed between detection and satisfactory mitigation of a cyber incident.
Purpose	Indicates the ability of the organization's incident process to effectively mitigate cyber incidents. Though depending on the severity of the incident, a lower value will lower the opportunity for an attacker to achieve his goals.
Differentiation options	– By incident type, e.g. DDoS vs. system intrusion vs. malware vs. phishing vs. ... – By service affected. A telecommunications provider might for instance quantify this metric separately for fixed (e.g. fibre) and mobile (e.g. 4G) data services. – By network segment affected, e.g. office network vs. production server infrastructure
Data sources	– Incident management processes and tools – Incident evaluation reports – Forensic investigation reports

M15. Adequacy of incident escalation

Definition	# cyber incidents that was unrightfully escalated or that the organisation failed to escalate
Purpose	Indicates the ability of the organization to assess the severity of a cyber incident and handle the incident at the correct escalation level. A higher number implicates lower ability to correctly assess cyber incidents.
Differentiation options	– By incident type, e.g. DDoS vs. targeted attack vs. malware vs. phishing vs. ... – By service affected (see examples under M14 and M15) – By network segment affected, e.g. office network vs. production server infrastructure
Data sources	– Incident management processes and tools – Incident evaluation reports – Crisis management procedures or logs/ documents

M16. Follow-up on forensic investigation

Definition	% forensic investigations that invoked tangible action (e.g. modify firewall rules or monitor on IoC) in operational security processes.
Purpose	Indicates the organisation's ability to translate outcome of forensic investigations into actual security enhancements. A higher percentage equals better performance.
Differentiation options	– By type of asset subjected to forensic investigation, e.g. workstation vs. server vs. network element vs. mobile device. – By entity that conducted the investigation, e.g. internal (CSIRT¹⁰, red team) vs. third party
Data sources	– Forensic investigation reports – Forensic investigation specialists

¹⁰ Computer Security Incident Response Team

M17. Effectiveness of forensic investigations

Definition	% forensic investigations where follow up action ultimately resulted in tangible effect (e.g. IoC sighting or blocking of traffic from illicit source)
Purpose	Indicates the organisation's ability to translate outcome of forensic investigations into actual security enhancements. A higher percentage equals better performance. Note: should be calculated as [# forensic investigations that resulted in tangible effect] / [# forensic investigations that invoked any follow up] (metric M16).
Differentiation options	– By type of asset subjected to forensic investigation, e.g. workstation vs. server vs. network element vs. mobile device. – By entity that conducted the investigation, e.g. internal (CSIRT, red team) vs. third party
Data sources	– Forensic investigation reports – Forensic investigation specialists – Security monitoring tools or SOC analyst (sightings) – Firewall logs or firewall maintenance staff – System logs or system operators

Ability to resist malware

M18. Follow up on cyber exercises

Definition	# cyber incident and/ or crisis exercises that was formally evaluated and for which improvement actions were formally defined.
Purpose	Indicates whether people involved are familiar with the cyber incident and management processes and if these processes are periodically improved. A higher number implicates higher probability that the processes are known and periodically improved.
Differentiation options	– By process that was exercised, e.g. cyber incident management, crisis management, SOC processes – By nature of exercise, e.g. process vs. technical (“capture the flag”) type exercises
Data sources	– Plans and schedules for cyber exercise – Evaluation reports of cyber exercises

M19. Adequacy of incident and crisis processes

Definition	Average # of major deficiencies revealed by cyber exercises.
Purpose	Indicates whether processes for incident handling and crisis management are sufficiently effective. A lower average number implicates higher effectiveness. Trend of the metric (score over time) reveals the organisation’s ability to learn and improve.
Differentiation options	– By process that was exercised, e.g. cyber incident management, crisis management, SOC processes – By nature of exercise, e.g. process vs. technical (“capture the flag”) type exercises
Data sources	Evaluation reports of cyber exercises.

M20. Effectiveness of malware prevention

Definition	# unique malware variants detected before activation
Purpose	Indicates the ability of the organization to prevent malware infections. The metric refers to “unique malware variants” because it aims to focus on: - technical prevention capability - if at least one instance of a particular malware sample is detected, this proves that the organisations has the technical means to do so, even if other instances were not detected as fast - threat level: malware samples with many occurrences are often mass phishing campaigns rather than targeted attacks against a particular organisation.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device. – By type of malware, e.g. rootkit vs. ransomware vs. adware vs. spyware vs. worm vs. ...
Data sources	– anti-virus logs – reputation filters – dynamic analysis (sandbox) detection

M21. Malware detection rate	
Definition	Monthly # of malware infections detected after activation, divided by monthly # of malware variants detected before activation (metric M20)
Purpose	This metric qualifies the organisation's detection capability. Only unique malware samples are included in the metric, as this filters out multiple occurrences of trivial mass phishing campaigns.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device. – By type of malware, e.g. rootkit vs. ransomware vs. adware vs. spyware vs. worm vs. ... Note: malware categories associated with mass campaigns should not be filtered out as these may be misused for targeted attacks as well.
Data sources	– IT support desk – Incident response register – anti-virus logs – reputation filters – dynamic analysis (sandbox) detection

M22. Timeliness of malware detection	
Definition	Mean # days between malware becoming active on the first system and detection of the infection
Purpose	This metric represents the organisation's monitoring and detection capability. Within the scope of targeted, multi-stage attacks, it makes more sense to consider the attack as a whole than to consider infections of individual systems. When the time of the first infection cannot be determined a best estimate should be used instead.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device. – By type of malware, e.g. rootkit vs. ransomware vs. adware vs. spyware vs. worm vs. ...
Data sources	– Forensic investigation reports – Security monitoring tools or SOC analyst

M23. Effectiveness of malware containment	
Definition	% systems/ assets infected by an identical malware variant before initial remediation
Purpose	This metric represents the organization's capability to contain malware infections. Initial remediation is defined as the action of cleaning up all infected systems and restoring them to their intended state. The term 'initial' reflects that later remediation actions may be required to clean up systems that were not initially identified as infected or not successfully disinfected. The metric can be influenced either by enhancing the detection and response capability or by impeding malware spreading.
Differentiation options	– By type of asset, e.g. workstation vs. server vs. network element vs. mobile device. – By type of malware, e.g. rootkit vs. ransomware vs. adware vs. spyware vs. worm vs. ...
Data sources	– IT support desk – Incident response register

M24. Effectiveness of malware remediation	
Definition	Mean # days between detecting a malware incident and finalising clean-up of all infected systems
Purpose	This metric represents the organisation's malware response capability, including the forensic activities to identify all affected systems.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device. – By type of malware, e.g. rootkit vs. ransomware vs. adware vs. spyware vs. worm vs. ...
Data sources	– IT support desk – Incident response register

Ability to resist system intrusions

M25. Effectiveness of intrusion prevention	
Definition	Yearly # of independent system intrusion incidents
Purpose	This metric qualifies a combination of prevention capability, detection capability and attack activity and thus serves as a basic indicator. Note: as a single targeted attack might compromise multiple systems/ assets, this metric counts independent (unrelated) intrusions only.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device. – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Incident response register – IDS logs

M26. Resistance against lateral movement	
Definition	Mean # systems compromised in a single attack
Purpose	This metric represents a combination of detection capability and attack activity and thus serves as a basic indicator. It also reflects the complexity and severity of targeted attacks encountered by the organisation.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Incident response register – Forensic investigation reports

M26. Resistance against lateral movement	
Definition	Mean # systems compromised in a single attack
Purpose	This metric represents a combination of detection capability and attack activity and thus serves as a basic indicator. It also reflects the complexity and severity of targeted attacks encountered by the organisation.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Incident response register – Forensic investigation reports

M27. Timeliness of intrusion detection	
Definition	Mean # days that elapsed between initial system intrusion and detection of the incident.
Purpose	This metric represents the organisation's monitoring and detection capability. Within the scope of targeted, multi-stage attacks, it makes more sense to consider the attack as a whole than to consider intrusions of individual systems.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Incident response register – Forensic investigation reports – Security monitoring tools or SOC analyst

Ability to resist DDoS attacks

M28. Effectiveness of intrusion remediation

Definition	Mean # days that elapsed between initial detection of a system intrusion and restoring security and normal operation of all affected systems.
Purpose	This metric represents the organisation's intrusion response capability, including the forensic activities required to identify all affected systems.
Differentiation options	– By type of asset targeted, e.g. workstation vs. server vs. network element vs. mobile device – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Incident response register – Forensic investigation reports

M29. Attack surface for system intrusions

Definition	% applications on (end user) system with no relevance for the end-user's tasks and responsibilities
Purpose	This metric gives an indication of the attack surface for system intrusions.
Differentiation options	– By type of asset, e.g. workstation vs. server vs. web portal – By asset criticality, e.g. vital vs. common, low-medium-high risk, internet facing vs. solely internal
Data sources	– Dedicated scan or investigation – Employee survey

M30. Timeliness of DDoS detection

Definition	Mean time (minutes, hours) required to acknowledge a DDoS attack, i.e. mean time elapsed between initial alert and formal diagnosis of an ongoing DDoS attack
Purpose	Indicates the organisation's ability to promptly recognize that it is enduring a (significant) DDoS attack. A low number equals better performance.
Differentiation options	– By magnitude or complexity of DDoS attack – By time of day/ week (office hours vs. nightly vs. weekend)
Data sources	– Logs from (self managed) anti-DDoS appliance – Reports from DDoS mitigation partner – Reports from Network Operations Center

M31. Service disruption due to DDoS attacks

Definition	# hours of service unavailability due to DDoS attacks
Purpose	Indicates the organisation's ability to continue its daily business and operations when enduring a (significant) DDoS attack. A lower number equals better performance.
Differentiation options	– By magnitude or complexity of DDoS attack – By time of day/ week (office hours vs. nightly vs. weekend) – By service affected (see examples under M14 and M15)
Data sources	– Reports from Network Operations Center

Ability to protect credentials

M32. Service degradation due to DDoS

Definition	# hours of service degradation due to DDoS attacks
Purpose	Indicates the organisation's ability to maintain quality of service in its business and operations when enduring a (significant) DDoS attack. A lower number equals better performance.
Differentiation options	– By magnitude or complexity of DDoS attack – By time of day/ week (office hours vs. nightly vs. weekend) – By service affected (see examples under M14 and M15)
Data sources	– Reports from Network Operations Center

M33. Operational impact of DDoS attacks

Definition	Mean time to restore operations after a DDoS attack, i.e. mean time elapsed between initial alert and resuming normal level of operations (corresponding to deactivation of DDoS mitigation)
Purpose	Indicates the organisation's ability to promptly eliminate an ongoing DDoS attack and the corresponding impact on its business and operations. A lower number equals better performance.
Differentiation options	– By magnitude or complexity of DDoS attack – By time of day/ week (office hours vs nightly vs weekend) – By service affected (see examples under M14 and M15)
Data sources	– Logs from (self managed) anti-DDoS appliance – Reports from DDoS mitigation partner – Reports from Network Operations Center

M34. Misuse of valid credentials

Definition	Annual # intrusion attempts that demonstrably involved unauthorised use of valid access credentials or tokens.
Purpose	Information on actual abuse incidents is a key indicator of the threat level. It should be kept in mind that it also reflects the organisation's capability to monitor access attempts for credential misuse.
Differentiation options	– By credential type, e.g. passwords vs tokens – By type of asset targeted, e.g. workstation vs. server vs. network element vs. database vs. critical application
Data sources	– System logs – IT support desk – Incident response register

M35. Timeliness of credential revocation

Definition	Mean time (hours, days) that elapsed between discovering loss or compromise of access credentials and revoking use.
Purpose	Stolen/captured credentials or credential generating tokens provide an entrance into the organisation's network only if the attacker is allowed the time window required to abuse them.
Differentiation options	– By credential type, e.g. passwords vs. tokens – By incident type, i.e. loss vs. demonstrable compromise
Data sources	– IT support desk – Incident response register

Ability to protect key assets

M36. Strength of user passwords

Definition	% user passwords that was successfully cracked during the most recent penetration test
Purpose	This metric provides an indication of security awareness of end users and of the current attack surface. It should be noted that even a relatively low score on this metric may represent a serious risk as cracking one user's password is often enough to obtain full access.
Differentiation options	– By category of password owner, e.g. end user vs. system operator vs. C-level executive – By account permissions level, e.g. root access vs. installation rights vs. read/write – By asset protected, e.g. workstation vs. server vs. network component vs. critical application
Data sources	– Penetration testing report

M37. Resistance to data exfiltration

Definition	Annual % data exfiltration attempts that was averted through automated or human intervention
Purpose	This metric represents a combination of detection capability and attack activity.
Differentiation options	– By type of sensitive data, e.g. personal vs. commercial vs. operational – By level of sensitivity (data classification), e.g. confidential vs. secret vs. mission critical – By nature of intervention, e.g. automated vs. human
Data sources	– DLP¹¹ logs – SIEM – Incident response register

¹¹ Data Leakage Prevention solution

M38. Presence of sensitive assets

Definition	# of assets for which a confidentiality or integrity breach is assessed to have high potential impact
Purpose	This metric reflects the organization's awareness about its key assets and the current attack surface. Here we note that attackers may be interested in a variety of IT assets, e.g. personal employee information that can be used to craft a social engineering attack to credit-card data, customer withdrawal limits that an attacker may attempt to modify and ATM machines that an attacker may aim to compromise.
Differentiation options	– By asset type, e.g. systems vs. applications vs. data – By nature of threat, e.g. confidentiality vs. integrity – By value for the attacker, e.g. end target vs. stepping stone (i.e. instrumental for getting to the end target)
Data sources	– CMDB – Risk register

M39. Accessibility of sensitive assets

Definition	% of sensitive assets for which access is granted on a need-to-know basis only
Purpose	This metric indicates the attack surface provided to attackers. Strict access control raises the bar for attackers to get to the desired asset or data.
Differentiation options	– By asset type, e.g. systems vs. applications vs. data – By value for an attacker, e.g. end target vs. stepping stone (i.e. instrumental for getting to the end target)
Data sources	– IAM¹² systems – Policies and procedures

¹² Identity & Access Management

M40. Confinement of sensitive data

Definition	Mean # systems on which a set of sensitive data is stored.
Purpose	This metric is an indication of attack surface. When copies of data assets are stored on different servers or locally, an attacker only needs to obtain access to the closest or least protected system.
Differentiation options	– By type of sensitive data, e.g. personal vs. commercial vs. operational – By level of sensitivity (data classification), e.g. confidential vs. secret vs. mission critical
Data sources	– CMDB – Configuration management systems – System operators

M41. Encryption of sensitive data assets

Definition	% of sensitive data assets that is structurally encrypted while stored on IT systems or transmitted over networks.
Purpose	This metric is an indication of attack surface. Cryptography, when applied well, shields attackers from sensitive data even when they have obtained access to the medium on which it is hosted. Note: refers to data that is encrypted by means of solid (strong) algorithms and key lengths.
Differentiation options	– By type of sensitive data, e.g. personal vs. commercial vs. operational – By level of sensitivity (data classification), e.g. confidential vs. secret vs. mission critical
Data sources	– CMDB – Configuration management systems – System operators

M42. Abuse of sensitive assets

Definition	Annual # incidents where a system/ application earmarked as sensitive was compromised
Purpose	This metric qualifies the organisation's ability to defend key assets. Note that it represents a combination of detection capability and attack activity.
Differentiation options	– By asset type, e.g. systems vs. applications vs. data – By value for the attacker, e.g. end target vs. stepping stone (i.e. instrumental for getting to the end target)
Data sources	– Incident response register – Forensic investigation reports

Ability to measure and minimize damage

M43. Monetary losses due to cyber incidents	
Definition	Money lost due to cyber incidents as a percentage of money transferred [i.e. total damage (cash out to attackers) in a certain period, divided by the amount of money transferred in that period]
Purpose	Indicates the direct monetary losses incurred by cyber incidents. A lower value indicates a better ability to prevent or reduce losses.
Differentiation options	– By market, e.g. retail vs. SME vs. corporate – By attack type, e.g. targeted attack vs. extortion vs.... – By platform or service targeted. A financial provider might for instance quantify this metric separately for ATMs, internet banking and debitcard transactions.
Data sources	– Incident response register – Financial reports – Product managers (through interviews)

M44. Penalties due to cyber incidents	
Definition	Absolute amount of money (in Euros) lost in penalties and/ legal liabilities as a direct and demonstrable result of cyber incidents.
Purpose	Indicates the penalties and legal costs to indemnify customers that have suffered (from the results of) cyber incidents (for instance, compensation for customers as a result of lost turnover because of failing payment systems). This metric also relates to penalties or fines that have to be paid to regulatory bodies as a result of not being compliant to cyber security regulations. A lower value indicates a better ability to prevent penalties or legal fines that result from cyber incidents and not being compliant to regulations.
Differentiation options	– By market, e.g. retail vs. SME vs. corporate – By attack type, e.g. targeted attack vs. extortion vs.... – By platform or service targeted. A financial provider might for instance quantify this metric separately for ATMs, internet banking and debitcard transactions.
Data sources	– Incident response register – Financial reports – Legal department (through interviews) – Product managers (through interviews)

M45. Customer loss due to cyber incidents	
Definition	# customers lost as a direct and demonstrable result of cyber incidents (churn)
Purpose	Indicates the ability of the organisation to maintain customer confidence following cyber incidents. A lower value indicates a better ability to prevent churn.
Differentiation options	– By market, e.g. retail vs. SME vs. corporate – By geographic region, e.g. NL vs. Europe vs. EMEA
Data sources	– CRM systems – Commercial teams (through interviews)

SRP Cyber Security

The Shared Research Program (SRP) Cyber Security is a joint R&D program in which TNO, ABN AMRO, Rabobank, ING and Achmea develop novel methods and technologies in the areas of cyber resilience, monitoring & response, cyber threat intelligence and secure transactions. The core purpose of these innovations is to enhance the prevention and detection of cyber attacks as well as the recovery thereafter. The SRP has a shared funding model that involves contributions from the SRP partners and the Dutch government. Project teams are comprised of specialists from all participating organizations and results are verified in a realistic setting (e.g. by running tests on authentic data collected in partner infrastructures). Much of the program's outcome is also shared with the broader cyber security community.

Interested parties from any industry are welcome to join the SRP Cyber Security. For more information please see <https://www.tno.nl/en/collaboration/partners-of-tno/shared-research-programme-cybersecurity/>.

M46. Reputational effects due to cyber incidents	
Definition	Factor increase in negative public statements as a result of cyber incidents [i.e. number of negative public statements in a certain period following a major cyber incident, divided by the average number of negative utterances in a similar timeframe before the incidents occurred]
Purpose	Indicates the possible loss of reputation as a result of cyber incidents. A lower value indicates a better ability to prevent damage to reputation.
Differentiation options	– By market, e.g. retail vs. SME vs. corporate – By geographic region, e.g. NL vs. Europe vs. EMEA
Data sources	– Social media (retail market only) – Customer survey (corporate market)

M47. Retrieval of financial losses	
Definition	% cyber incident related losses that has been retrieved, e.g. by legal actions or cyber insurance.
Purpose	Indicates the ability of an organization to reclaim losses (due to cyber incidents) that it suffered after the incident has been mitigated, thereby in total minimizing the losses.
Differentiation options	– By market, e.g. retail vs. SME vs. corporate – By attack type, e.g. targeted attack vs. extortion vs.... – By platform or service targeted. A financial provider might for instance quantify this metric separately for ATMs, internet banking and debitcard transactions.
Data sources	– Financial reports – Financial department (interviews) – Legal department (interviews)

References

[Lockheed] E.M. Hutchins et al., Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains, Lockheed Martin Corporation, 2011.
<http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf>

[Measuring] Richard Kerkdijk and Paul Samwel, Measuring cyber resilience, Innovating in Cyber Security pp 9-14, May 2017
<https://www.tno.nl/media/9419/innovating-in-cyber-security.pdf>

Further reading:

A. Jacquith, Security Metrics – Replacing Fear, Uncertainty and Doubt, Addison-Wesley, 2007

W. Krag Brotby and Gary Hinson, Pragmatic Security Metrics - Applying Metametrics to Information Security, CRC Press, 2013



