Targets

#	Country
1	United States
2	Germany
3	Netherlands
4	United Kingdom
5	Switzerland
6	Belgium

Origins

#	Country
1	China
2	Russia
3	United States
4	Ukraine
5	India
6	Thailand

Innovating in Cyber Security

Shared research 2017

Rabobank
ABN·AMRO
ING BANK
TNO innovation for life

Contents

Preface

This magazine is the first edition of the Dutch Cyber Security Shared Research Program (SRP), a collaboration involving TNO, ING, ABN AMRO, Rabobank and Achmea. The magazine serves two purposes; to share our experiences in cooperating in a shared research Program and to share some of the results that have been achieved in this Program. We hope these experiences and results offer you some fresh perspectives on cybersecurity innovation, which we believe is essential to maintain a prosperous society.

We strongly believe in the value of cooperation, which is the basis for this Program and is a common theme that returns in the articles in the magazine. Throughout the magazine, leaders from each of the partners involved in this SRP share their views on cooperation within the Program, and on the resultant benefits.

We will continue the cooperation and foster new partnerships in the coming years.
We encourage you to build on experiences and results that we present. We trust that this magazine will inspire you to innovate in cyber security and cyber resilience and explore new ways of improving your defence against cyber attacks.

Enjoy reading the magazine!

Mark Wiggerman (ABN AMRO)
Tom Huitema (Achmea)
Tommy Koens (ING)
Paul Samwel (Rabobank)
Reinder Wolthuis (TNO)

Cyber Security Shared Research Program

The Cyber Security Shared Research Program (SRP) is a unique research and innovation Program. It provides a context in which partners can cooperate to improve cyber security by means of innovation in various technologies and processes. The SRP involves a year-long commitment, in which the partners explore four lines of research:

- Monitoring & Response – the aim is to improve the detection of cyber security incidents, and improve the response to incidents once they are detected, through innovation in monitoring and response technologies and processes.
- Controlled Resilience – the aim is to improve organizations' cyber resilience, through innovation in resilience technologies and processes. Cyber resilience is defined as an organization's ability to cope with cyber attacks on its infrastructure or electronic services.
- Cyber Intelligence – the aim is to share threat intelligence more effectively, and to use it for the early detection and prevention of cyber attacks.
- Secure Transactions 2.0 – the aim is to define and design a security architecture for the generic transaction method or platform of the future; this architecture must be independent of specific technology or channels, and its potential range of applications should preferably not be restricted to the financial sector.

The parties currently involved in the Cyber Security SRP are TNO (the Netherlands Organization for Applied Scientific Research) and various Dutch financial institutions (ABN AMRO, ING, Rabobank and Achmea). Interested parties from any sector are welcome to join the SRP.

The goal of the Cyber Security SRP is to improve the prevention and detection of cyber attacks (and the subsequent recovery) by developing a range of innovative technologies and methods.

The SRP provides a context in which the partners can cooperate to improve cyber security by means of innovation in technologies and processes

This development work will draw on the participants' expertise in the areas of security technologies and methodologies, data analytics, incident and crisis management, and behavioural sciences.

The development of innovative technologies and methods will benefit the Program's partners by:

- improving their ability to control cyber security risks;
- further enhancing the maturity of cyber security in the financial sector.

This will help to reduce the losses caused by cyber attacks and to increase customers' confidence in the security of digital services.

The SRP participants are seeking cooperation with universities, vendors and government agencies

The CISO Advisory Board: Martijn Dekker (CISO ABN AMRO), Wim Hafkamp (CISO Rabobank), Reinder Wolthuis (SRP program manager TNO), Vincent Thiele (manager CCERT ING), Henk Jan Vink (Director Networked Information TNO), Dimitri Hehanussa (Business development TNO)

Advantages of cooperation

The SRP focuses on three areas of cooperation:

- Shared workload – while the program's project teams are primarily made up of TNO staff, these are complemented with staff members from each of the participating partners.
- Shared data – the participating partners provide anonymized, real-life data to evaluate innovative security methods.
- Shared funding – each partner pays part of the costs of the Program, in addition to various practical contributions. The Dutch government also provides funding.

In the projects, the SRP participants are also seeking cooperation with universities, vendors and government agencies such as the Dutch National Cyber Security Centre (NCSC).

The research results delivered by this unique collaboration are far more effective than anything the individual partners could achieve by themselves. TNO's research-oriented approach blends with the more practice-oriented, operational approach adopted by the other SRP partners. The results can be directly verified using real-life (anonymized) data. Occasionally, research results can be implemented directly into the SRP partners' infrastructure.

The SRP and cyber security research in general

Cyber security research and development is undergoing rapid development in the Netherlands (see Figure 1) and elsewhere in the world.

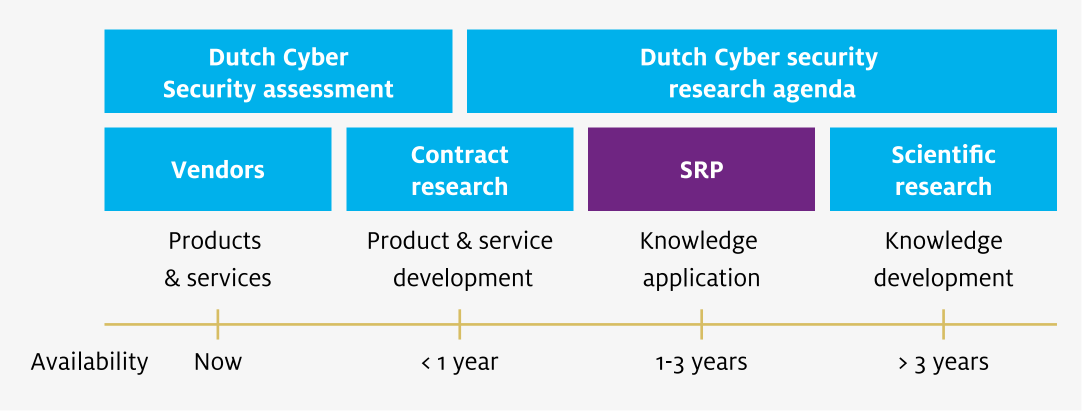

Dutch Cyber Security assessment

Dutch Cyber security research agenda

Vendors

Contract research

SRP

Scientific research

Products & services

Product & service development

Knowledge application

Knowledge development

Availability

Now

< 1 year

1-3 years

> 3 years

Figure 1: SRP program in the context of other cyber security research

On the right side of the ‘spectrum’ shown in the figure above, scientific research produces scientific knowledge. Such knowledge usually takes time to develop, tends to be rather elementary and cannot be applied directly. On the left side of the ‘spectrum’, however, we see vendors implementing solutions in their products. These short-term solutions are market ready. Before such solutions can be implemented, product & service development is needed. The majority of research activities in the Cyber Security SRP focus on ‘knowledge application’. This means that the SRP’s research activities provide knowledge that can be used as input for product & service development. In a minority of projects, the research activities involved are more oriented towards long-term goals.

Structure and roadmap

The Cyber Security SRP is headed by a TNO program Manager and governed by the CISO Advisory Board, which is made up of the various banks’ Chief Information Security Officers (CISOs). TNO and the CISO advisory board are jointly responsible for roadmap development. Details of the current roadmap are shown in Figure 2 below. During the course of the year, the progress and output of each project is closely monitored. Each project is classified into one of three categories:

1. new, applied research; the SRP partners have little or no knowledge of the subject, there are few if any commercial products, and no de-facto standards are available;
2. the SRP partners have a limited but developing knowledge of the subject; the portfolio of commercial products is growing, but there are still major developments going on, (de-facto) standards are emerging;
3. a great deal of research has been done on the subject, both in the context of the SRP projects and elsewhere; as a result, our partners have quite a good knowledge of the subject; Common Of The Shelf (COTS) products are becoming widely available and most of them comply with available de-facto standards.

Each program participant is represented at an annual workshop, where potential new lines of research are discussed

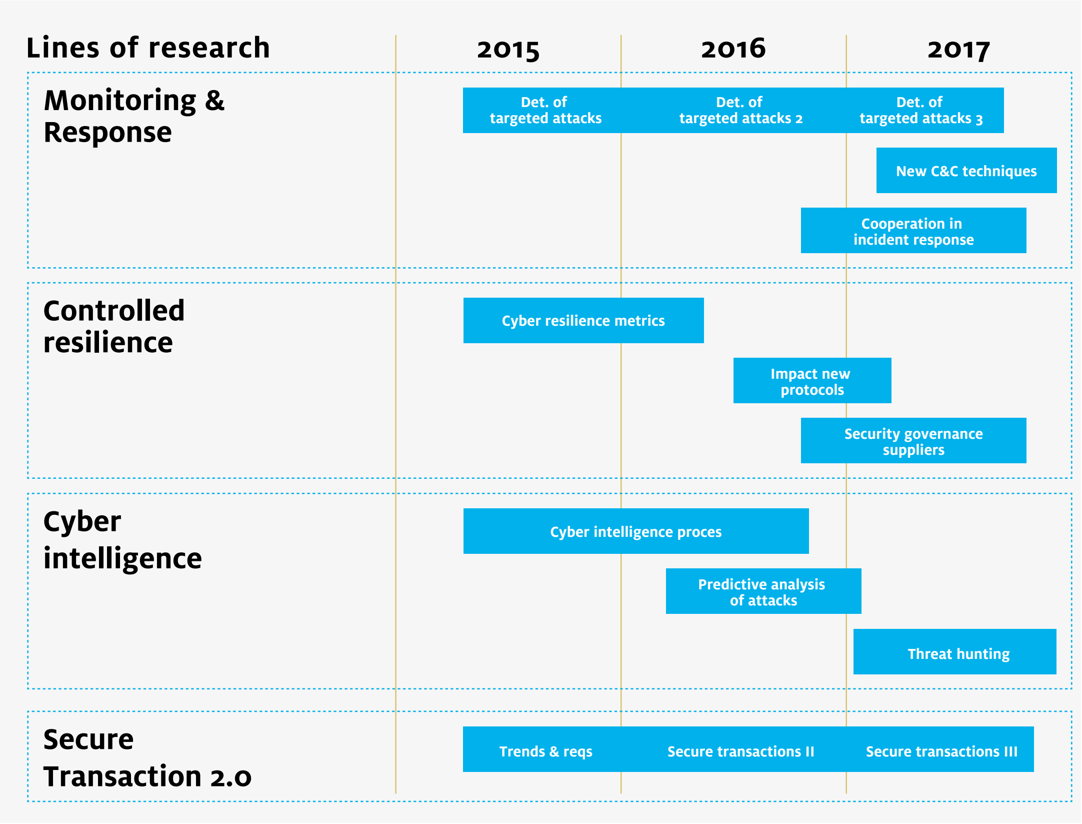

Figure 2: current SRP roadmap

Our goal is to continue to pursue research in the first two categories, as long as the subject is of interest to the SRP partners. Any subjects that are introduced into Category 3 will be passed on for follow-up outside the framework of the SRP. The focus will then be on technology transfer and on commercializing the results.

Results

The SRP delivers many different types of results, ranging from methodologies to frameworks, algorithms, and proof of concepts, for example. Although the Intellectual Property Rights (IPR) arising from the SRP are owned by TNO, every SRP partner has unlimited rights regarding their use. The SRP's goal is to enable as many parties as possible (even those outside the SRP) to benefit from the results. A range of technology transfer approaches are used for this purpose:

- publishing results and encouraging other parties to benefit from them;
- cooperating with vendors to integrate new solutions into their products;
- launching spin-off companies to integrate the results into new products;
- selling the IPR to interested parties and investing the earnings in new SRP research;
- allowing others to use the results under a creative commons licence; third parties can distribute, share or, in some cases, also process the content covered by the IPR.

This was the first cooperation of its kind, and as such it took some time to develop a suitable format to carry out the activities involved. Now, however, the first beneficial results are starting to emerge.

SRP's goal is to enable as many parties as possible (even those outside the SRP) to benefit from the results

“ Participating in the Shared Research Program is a great way to connect top talent from different organizations in a setting where they can learn, experiment and share. In addition to generating quality output, the SRP triggers creativity.
Also this joint research helps the security professionals to get to the next level in cybercrime thinking and to become more effective. I am convinced that cooperative ventures like this will enable us to continue to provide practical, manageable security for banking in cyberspace. This will not only be good for ABN AMRO, it will benefit society as a whole.”

Martijn Dekker
CISO ABN AMRO

Measuring cyber resilience

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 taken steps to 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)?

These and other compelling questions evoked a strong desire among financial institutions to measure and quantify cyber resilience within their organizations. Thus a project was launched to jointly define a meaningful framework of cyber resilience metrics. This article presents the framework's structure and underlying philosophy, as well as examples of the actual metrics. It also describes some of the lessons learned.

Point of departure

In itself, the concept of security metrics is not entirely new. Numerous articles have been published[1] on the subject and many organizations – including the Dutch financial institutions involved in this initiative – already apply them in some way, shape or form. Existing structures of security metrics, however, are characterised by some distinct limitations:

- The selection of metrics was often driven by ease of implementation. More often than not, security status reports are largely derived from performance dashboards that are readily available in technical security solutions (e.g. anti-virus tooling or IAM[2] platforms). While virus counts and failed login attempts are useful parameters to consider, it is unlikely that they truly meet the information needs of strategic stakeholders such as the CISO or executive leadership[3].
- Rather than reflecting actual effects or performance, metrics for security often 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 completed (mandatory) security training[4]. In itself, however, the fact that an employee has completed a given e-learning module offers little assurance that he or she will exhibit appropriate behaviour when faced with an actual security threat (e.g. a phishing email).

[1] A selection of relevant literature is included in the backof this publication.

[2] Identity & Access Management

[3] On top of this, many such (technical) metrics are heavily affected by the current level of inherent threats. The number of malware infections intercepted at an organization's network gateway, for instance, depends on the strength of its defensive capabilities but also on the actual number of attempted attacks during the period under consideration.

[4] This limitation of traditional security metrics also came to light in a recent course presented by Delft University of Technology, which indicates that this issue is not exclusive to the financial industry.

The second issue is often a reflection of security cultures that are driven by compliance objectives. Until recently, such cultures also prevailed in the financial industry and this has greatly influenced the nature of the metrics and reporting formats for (cyber) security that are presently in use.

Security metrics with the above characteristics offer limited insight into the actual status and performance of cyber resilience measures. Correspondingly, they are not particularly suitable for managing security operations or justifying investments. This initiative pursued a material step forward by focusing on metrics that reflect an organization's cyber resilience abilities and the actual effects achieved through technical and organizational security measures (see Figure 1).

Building a meaningful framework

In order to establish meaningful metrics, they must be based on a common understanding of the term 'cyber resilience'. To this end, the following definition was used:

> Cyber resilience is the ability of an ecosystem (e.g. an organization, infrastructure, system) to
> ...withstand deliberate attacks on technical infrastructure that are conducted from cyberspace
> ...rapidly recover from the adverse effects of such attacks
> ...limit the damage of such attacks on business, people and society
> ...prepare for and adapt to changing conditions[5]

This definition confirms that metrics for cyber resilience should reflect an organization's abilities and performance rather than the specific controls or actions that it has put in place. It also indicates that a set of metrics for cyber resilience will only be meaningful if it covers the full life-cycle of preventing, detecting and responding to 'deliberate attacks'. While such attacks can be diverse in nature (e.g. in terms of sophistication and underlying motivation), this project specifically focused on so called 'targeted cyber attacks', also referred to as 'Advanced Persistent Threats' (APTs).

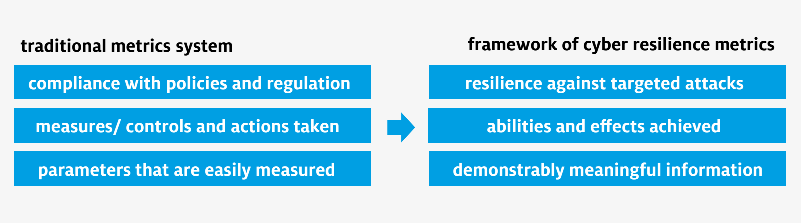
traditional metrics system
compliance with policies and regulation
measures/ controls and actions taken
parameters that are easily measured
framework of cyber resilience metrics
resilience against targeted attacks
abilities and effects achieved
demonstrably meaningful information

Figure 1: Enhancing the intrinsic nature of security metrics

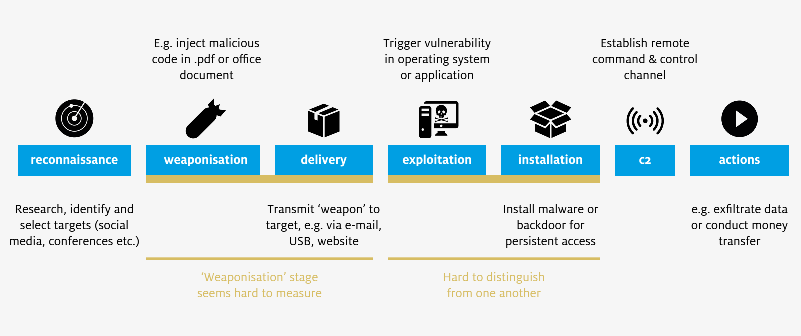
E.g. inject malicious code in .pdf or office document
Trigger vulnerability in operating system or application
Establish remote command & control channel
reconnaissance
weaponisation
delivery
exploitation
installation
c2
actions
Research, identify and select targets (social media, conferences etc.)
Transmit 'weapon' to target, e.g. via e-mail, USB, website
Install malware or backdoor for persistent access
e.g. exfiltrate data or conduct money transfer
'Weaponisation' stage seems hard to measure
Hard to distinguish from one another

Figure 2: Cyber kill chain and merger of stages for metrics framework

This class of attacks is of particular concern to financial institutions, as historic cases have proven that they are a very real threat with a high potential impact.

A widely recognized model for targeted cyber attacks is the 'cyber kill chain' developed by Lockheed Martin. The overarching structure for the framework of cyber resilience metrics was based on this model, albeit that some stages were merged as the distinction between them was not considered useful in this context (see Figure 2).

To incorporate meaningful metrics into this structure, 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.

The exercise produced a total of 47 metrics that correspond to relevant cyber resilience abilities across the various stages of the cyber kill chain. These metrics were consolidated into 10 core categories:

[5] e.g. changes in the methods used by attackers or in the organization's IT infrastructure

The fact that an employee has completed a given e-learning module offers little assurance that he or she will exhibit appropriate behaviour when faced with an actual security threat

1. Avert social engineering. This category 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).
2. Engage threat intelligence. Metrics in this category reflect the organization's ability to anticipate imminent or emerging threats before an actual incident occurs (e.g. a network intrusion). This is typically achieved by collecting and processing threat information, e.g. concerning new methods used by attackers[6].
3. Address vulnerabilities. Cyber attacks usually involve the exploitation of vulnerabilities in networks, systems or software. This category reflects the organization's ability to proactively discover and remediate such vulnerabilities.
4. Handle cyber incidents. There is a widespread notion that security incidents cannot be completely avoided if the adversary is sufficiently motivated and competent. Thus an organization's resilience relies heavily on its ability to detect and mitigate such incidents. This category reflects the status of that ability.
5. Resist malware. This category reflects the organization's ability to detect, contain and remediate malicious software that is present and/or active within its technical infrastructure. Typical attack scenarios involve the use of malware in various stages of the cyber kill chain.
6. Resist system intrusions. Metrics in this category reflect the organization's ability to defend against system intrusions and against factors such as 'lateral movement'[7] that are typical of targeted cyber attacks.
7. Resist DDoS attacks. This is a special category, since it is not directly associated with APT-type attack scenarios. The ability to detect and respond to Distributed Denial of Service (DDoS) attacks was included because it is a related issue that is of great concern[8] to large organizations.
8. Protect credentials. Targeted attacks often involve the abuse of credentials (e.g. user names and passwords) to gain access to systems or data. This category of metrics reflects the organization's ability to minimize such abuse, e.g. by quickly revoking any credentials that have been compromised.
9. Protect key assets. This category reflects the organization's ability to shield its most valuable system and information assets from abuse. Such key assets are typically the ultimate target of an APT-type attack.

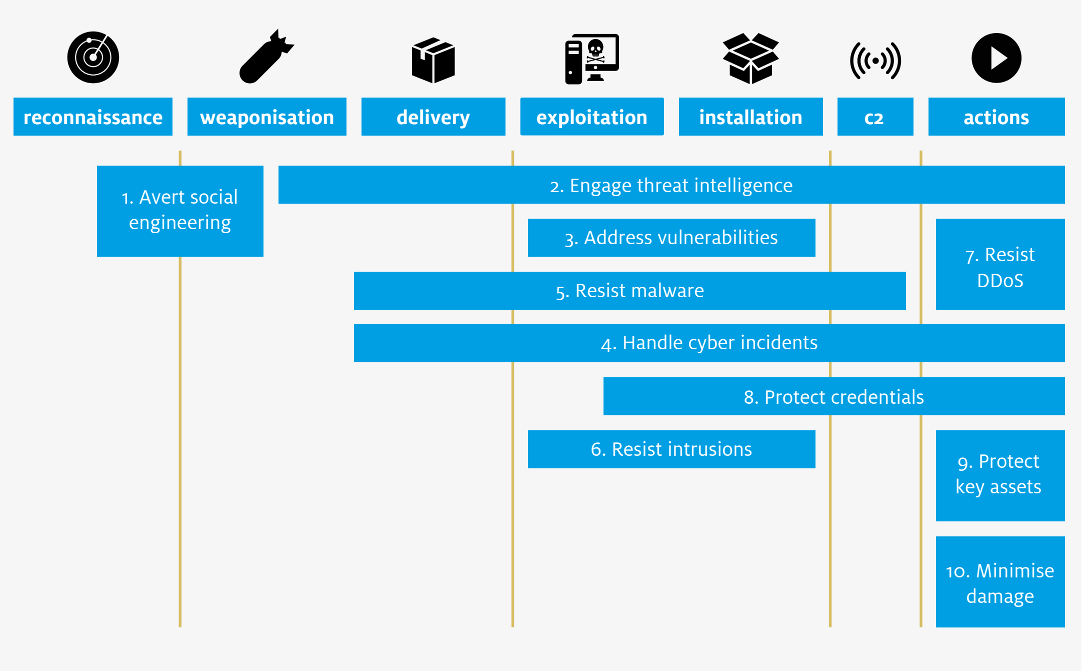

Figure 3: Core categories of cyber resilience metrics

[6] The concept of Cyber Threat Intelligence (CTI) is addressed at greater length in the article 'Towards a mature CTI practice', also included in this publication.

[7] Techniques used by cyber adversaries to progressively move through a network in search of the data, assets or resources that are ultimately targeted.

[8] This concern derives from developments such as the volumetric rise of DDoS attacks.

This involves factors such as resistance to data exfiltration and the accessibility of sensitive assets.

10. Measure and minimize damage. Cyber attacks can involve various types of damage, e.g. monetary or concerning the organization's reputation. This category of metrics reflects the organization's ability to control such damage.

The top-level framework structure is depicted in Figure 3.

As the figure shows, the core categories cover the full extent of the cyber kill chain, albeit with emphasis on later attack stages. It also reveals that numerous factors (many of which may not yet be actively monitored) must be taken into consideration during the complex process of measuring an organization's cyber resilience.

Metrics library

The metrics for cyber resilience were captured in a comprehensive library. Individual organizations can select the specific elements they wish to measure and monitor. An example from the 'ability to avert social engineering' category is shown below[9].

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

> To achieve the maximum effect, the information needs of individual stakeholders need to be matched with specific elements of the metrics framework

This example is clearly a product of the *measure abilities and effects* philosophy. It focuses on the actual behaviour that employees exhibit when exposed to phishing (rather than the completion of mandatory e-learning modules, see above).
It also reveals that measuring cyber resilience can be something of an undertaking. In this particular case, conducting a dedicated exposure test (i.e. simulating a phishing campaign) seems to be the most (if not the only) effective approach to collect the desired measurement data.
This may involve a substantial investment of time and money.

Notably, the framework also includes metrics that are much easier to implement. This is illustrated by the following examples, both from the 'ability to address vulnerabilities' category (see above).

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).

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.

These metrics require source data that should (at least for the most part) be readily available in vulnerability management tools and penetration testing reports (where such testing has taken place). Thus, some metrics in the framework can be adopted without requiring material investments.

It should be noted that the metrics included in the framework vary in terms of their nature and level of detail. Some indicate the overall state of a

[9] As explained above, each category is comprised of several metrics that jointly reveal the state of a cyber resilience ability. The examples described in this article generally focus on a specific performance aspect.

specific ability, whereas others focus on particular details. This is illustrated by the following examples:

M34. Misuse of valid credentials	
Definition	Annual # intrusion attempts that demonstrably involved unauthorised use of valid access credentials or tokens.

M35. Timeliness of credential revocation	
Definition	Mean time (hours, days) that elapsed between discovering loss or compromise of access credentials and revoking use.

The first is a top-level indicator of the organization's 'ability to protect credentials' (Category 8, see above), whereas the second focuses on a specific factor ('root cause') that may result in weak or strong performance in this area. Both are valuable in their own right, but typically for different target audiences. We will return to this in the following section.

Lessons learned

The major lessons learned from jointly developing a framework of cyber resilience metrics and implementing these in individual organizations were as follows:

1. Although it is indeed worthwhile to create effect-oriented metrics that reflect an organization's cyber resilience capabilities, such metrics can be hard to measure in actual practice. In most areas, it was possible to define metrics that reveal residual risk and the actual effects of cyber security measures. The following example is a product of this philosophy:

M21. Malware detection rate	
Definition	Monthly # of malware infections detected after activation, divided by monthly # of malware variants detected before activation.

This metric enhances an organization's assessment of malware resilience by correcting actual malware hits for generic increases in malware variants. The number of infections detected after activation can (for instance) be measured by counting malware related desktop/laptop re-enrolments. In this case, the metric is feasible because the underlying threat is covered by a measurable second line of defence. Such a second line is, however, not always in place and if it is insufficiently effective this might have an adverse effect on the metric.

2. Stakeholders are rarely interested in the full set of cyber resilience metrics. To a great extent, the information needs of individual stakeholders (e.g. security coordinators, technical specialists, business managers, senior leadership) depend on their role and position in the organization. Thus the value of a given metric is often a matter of perspective. Team leaders in a Security Operations Centre (SOC), for instance, are likely to benefit from metrics that address incident handling (Category 4, see above). However, they might be less inclined to measure the organization's readiness for social engineering schemes (Category 1). Similarly, CISOs or risk managers might focus on metrics that reflect the overall state of cyber resilience (capabilities) whereas those who coordinate operational (e.g. security) processes might prefer more specific assessments. To achieve the maximum effect, the information needs of individual stakeholders need to be matched with specific elements of the metrics framework.

3. Addressing the full set of cyber resilience metrics is challenging and requires a focused effort. Many of the cyber resilience metrics defined in this project require source data that is not readily available (or at least not actively collected) in the organisation or infrastructure of financial providers. To structurally quantify such metrics, changes must be made in existing system configurations, working procedures and reporting formats.Acquiring all of the data required for every metric defined in the framework will usually be too ambitious. Instead, it is more realistic to start with one or two metrics in each of the core categories (see Figure 3) and then expand the range according to specific needs (e.g. to

acquire any supplementary insights required by key stakeholders).

4. Comparing actual cyber resilience measurements across organisations requires a level of alignment that is presently not in place. The initial ambition was to compare the outcomes of measurements between the financial institutions involved. However, there is little overlap between the metrics selected by individual participants. Moreover, in areas where participants do use the same metric (e.g. malware and phishing losses in internet banking), the measurements are greatly affected by individual definitions and implementation choices. This is often due to differences in the operational resources from which source data can be collected (e.g. technical security solutions that may or may not be in place, or that may be produced by different vendors). The net result is that the outcomes of metrics measurement are not always comparable. So, instead of comparing their measurements with those of other organizations, the financial institutions involved will initially compare sets of their own, internal measurements taken at different times.

In view of the third point, the parties involved will not – at least for the time being – pursue a unified, industry-wide cyber resilience standard. It is, however, conceivable that there will be a renewed interest in (and even a requirement for) a normative standard at some point in the future. With this in mind, every metric incorporated into the framework allows the definition of quantifiable target levels to be defined at a later stage of development.

The way forward

The financial institutions involved in this activity will need some time to experiment with the concept of cyber resilience metrics. Indeed, they will need several quarters' worth of data and practical experience to effectively assess the practical value of many of these metrics. It might be possible to quantify some metrics using historical data, which would accelerate the process to some extent. In most cases, however, the required source data had either not been recorded or there was previously no need to retain it.

On the whole, the process of defining a framework for cyber resilience metrics has provided the participants with a useful reference point for building or enhancing their own internal dashboards. The value of this experience is likely to increase further when the partners exchange details of their initial practical experiences (e.g. with regard to the collection of internal source data or use of these resilience metrics in actual decision-making).

“One aspect of the Shared Research Program that really fires my enthusiasm is the way that each project stimulates creativity and out-of-the-box thinking. I sincerely hope that more organizations will follow this Dutch banking initiative, and contribute to the learning curve in this way.”

Wim Hafkamp
CISO Rabobank

The future of secure transactions

The continuous evolution of the transaction web and its supporting technology compels us to update our assessment of cyber security in the area of transaction services. We foresee that future transaction services will require new ways of assuring trust and new ways to mitigate fraud. There will also be a need for the effective validation of new types of relations, especially in the consumer-to-consumer and the device-to-business segments. In this context, we envision a need for 'conditional transactions'. These will allow users to define customizable and automatically verifiable conditions that must be fulfilled before a transaction can be executed. This paper summarizes current trends and future transaction services. It also presents a blueprint for the proposed Secure Transaction 2.0 (ST2.0) ecosystem, together with an analysis of the cyber security challenges involved.

The emerging Internet of Things (IoT) is launching a new category of players – devices – onto the transaction web

Trends

One of the principles behind Bank2.0 and Bank3.0 is that future financial services will be driven by customer behaviour: customers will choose 'those channels and interactions that get them to their desired solution in the quickest, most efficient manner'. One aspect of this trend is the rapid adoption of electronic and mobile payments, resulting in the move towards a 'cashless world'[22]. Customers are demanding greater ease-of-use, and their preferences are shifting and changing faster than ever before. As a result, ST2.0 services will need to be increasingly customer-oriented (as opposed to being procedure-driven) and customizable.

The most prominent technological trends in the financial sector right now are the rise of blockchain-based applications (including cryptocurrencies) and the introduction of innovative services by FinTech companies. One reason for the attractiveness of cryptocurrencies is that the transaction history is resilient to unwanted changes, another is that it is publicly available in a distributed ledger. No bank or government has control over the currency, nor are the transactions directly linked to a user's personal data. The rise of blockchain technologies has inspired the creation of new services based on the same concepts. Ethereum, for example, extends the scope of blockchains beyond the realm of currencies. They can be used to draw up smart contracts, where users 'can create their own arbitrary rules for ownership, transaction formats and state transition functions'[23].

These prominent emerging payment rails include peer-to-peer networks and mobile value transfer networks, which usually rely on a trusted intermediary third party to transfer value (mostly in small amounts) rapidly from one user to another. This ST2.0 ecosystem compels us to review the demand for global, transparent, cheap and ever faster transaction processing, and to reconsider the cyber security measures used in traditional processing systems. For example, the authorization procedures and fraud detection verifications involved in current cyber security measures may be too time-consuming in a world where faster processing is the norm.

Furthermore, the emerging Internet of Things (IoT) is launching a new category of players – devices – onto the transaction web. Pilot tests have already been carried out with washing machines and heating equipment, equipped with sensors that initiate maintenance requests over the internet[24]. The next logical step is to perform the corresponding financial transactions. With IoT devices projected to number 50 billion by 2030[25], machine-initiated transactions may become a large part of all future ST2.0 transactions. The ST2.0 ecosystem will therefore need to be able to support the new 'IoT transaction channel'. It will also need to tackle the cyber security challenges associated with increasingly autonomous devices. The above trends will lead to a restructuring of business roles in the transaction web. The upcoming financial PSD2 regulation (Revised Directive on Payment Services[26], which will require traditional payment service providers to

[22] The Future of Financial Services, World Economic Forum, June 2015, www3.weforum.org/docs/WEF_The_Future_of_Financial_Services.pdf

[23] https://github.com/ethereum/wiki/wiki/White-Paper

[24] The Internet of Things and payments, Part 2: The ghost of payments present, http://bit.ly/2c7FdY7

[25] https://451research.com/report-short?entityId=89399

[26] http://bitscan.com/articles/permissioned-and-unpermissioned-blockchains-part-1

open their infrastructures to other parties) will also galvanize the unbundling of business roles in the transaction web. In addition, linking-up transaction systems will create opportunities for delivering new services or for improving existing ones. For example, track & trace events in a service delivery transaction (e.g. 'product delivered') could be used to automatically initiate a corresponding payment transaction or vice versa. Indeed, some of the emerging payment rails already provide programmable interfaces (e.g. PayPal's API) to facilitate this coupling of transaction systems. The upcoming unbundling of transaction chains will introduce new interfaces between transactions systems. These will need to be secured, and will require new ways of ensuring end-to-end security.

Future transaction services

ST2.0-related trends will drive the development of innovative transaction services. Three such services (which could be available to consumers in the next five years) are described below.

Machine-initiated transactions may become a large part of all future ST2.0 transactions

'I owe you' – IOU

As presently conceived, IOU transaction services will be based on an 'IOU group'. This group will keep track of the various participants and their individual transactions. IOU transaction services make it possible to combine different ST2.0 trends, such as the use of emerging payment rails and customer empowerment. However, IOU-type services pose new cyber security challenges. For example, how can trust be assured in a peer-to-peer IOU group? Then there are 'conditional transactions', i.e. transactions that require the verification of a clearly specified condition (e.g. when one member is required to approve a purchase on behalf of the group). How can such conditional transactions be enabled before the transaction is executed? How can fraud be prevented in a peer-to-peer IOU ecosystem?

Smart domestic appliances

At some point in the evolution of the Internet of Things and Smart Homes, domestic appliances (such as refrigerators, washing machines, heaters and TVs) can be expected to initiate transaction services based on predefined rules set by their owners. Such developments will also be driven by trends such as customer empowerment (in terms of customizing the level of machine autonomy) and new third-party roles for equipment manufacturers. In this context, emerging payment rails (such as blockchains) might be used to verify the need for a conditional approval from the user and to issue such approval. The prospect of machines initiating financial transactions raises new cyber security questions. Our study identified the following specific issues. How can these new relationships between IoT devices and their owners be validated? How can IoT devices' authorization permissions and their owners' proof-of-consent be made transparent and traceable, while preserving the confidentiality of such information? Could a standardized protocol be introduced to counter a tendency towards ever more device-specific and/or vendor-specific solutions?

The prospect of machines initiating financial transactions raises new cyber security questions

Smart contracts

The goal is for ST2.0-related trends to enhance current transaction processes, or to make them more convenient. Take, for example, the process of creating and executing construction depot transactions.

When renovating their house, or building a new one, consumers can use a 'construction depot' to pay the builder/renovator in question. This means that, with the aid of a third party (e.g. a bank, notary or 'construction director'), the consumer places the total contracted amount for the renovation or construction work on deposit. The third party then makes sure that the builder/renovator receives the agreed portion of the total amount once a certain construction milestone is completed. To this end, a contract must be prepared and agreed between the consumer and the builder/renovator, setting out the conditions of fulfilment for each payment milestone.
The process of verifying the conditions is usually handled by the third party. This entire process involves considerable administrative work by the third party. If the construction depot was subject to a self-verifying, blockchain-based smart contract, this could greatly reduce the third party's workload, to the point where they would only need to provide approval for certain conditions or intervene in the event of a dispute. A construction depot based on a smart contract would involve the following cyber security challenges:

- trust between the buyer and the constructor/renovator;
- the trust of both parties in the ability of smart contracts to enable a construction depot transaction;
- confidentiality and integrity of the content of the smart contract (i.e. only authorized parties should be allowed to read it or modify it).

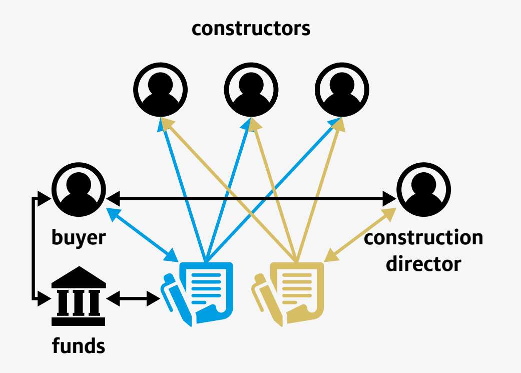

Figure 1: Relations in a construction depot based on a smart contract.

Blueprint

The use cases illustrate the potential for future, conditional transaction services. One way to greatly improve the verification of conditional transactions that involve many different stakeholders and sources of information is to make them reusable and customizable, by means of smart contract technology, for example. In this connection, we have identified the need for a shared ST2.0 service platform and for standard interfaces between the various systems making up this platform. Based on experiments we conducted with Ethereum, and on the cyber security challenges identified in the use cases, we have prepared a blueprint for the future ST2.0 platform (see Figure 2).

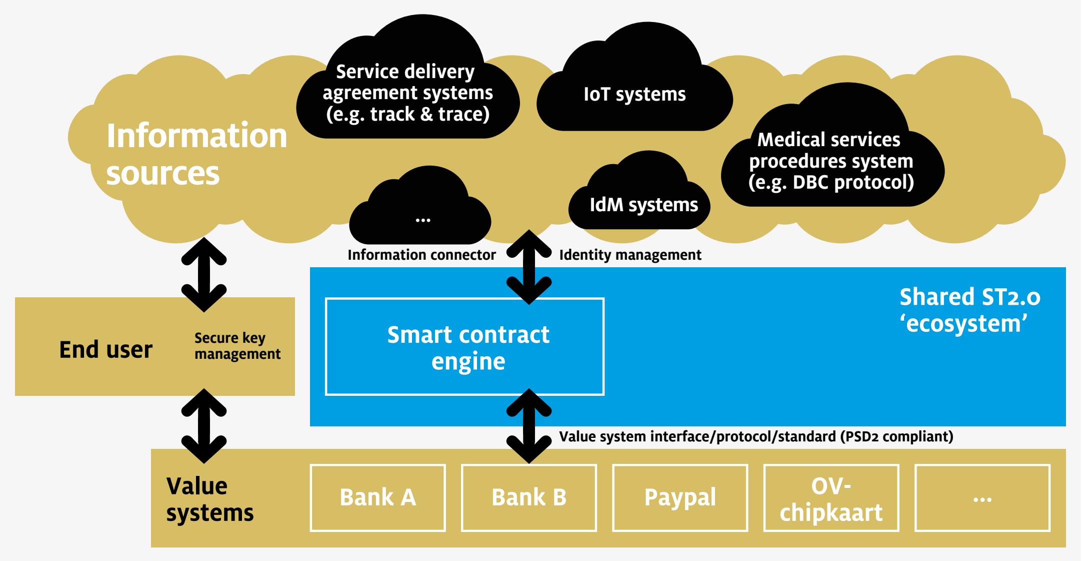

Figure 2: Secure transactions 2.0 blueprint

The shared ST2.0 platform will interconnect (by means of yet-to-be specified standard interfaces) with:

- Information sources. This information will be used to trigger the exchange of value. Accordingly, it needs to be reliable and someone needs to vouch for its validity. This may lead to the introduction of new roles and may also require additional Identity Management infrastructure and agreements (including legal agreements).
- Value systems. Smart contracts will require trustworthy, secure and transparent interfaces (supported by standards) with these value systems, which may even be connected indirectly through various 'hops'.
- End users. One important issue is whether end users interact with the envisioned ST2.0 ecosystem directly or via intermediaries. It is also important to ensure that such automated smart contracts accurately convey the users' intent.

Cyber security challenges

ST2.0 platform

Before the ST2.0 blueprint can be developed any further, various architectural and cyber security design issues need to be clarified. We will use the construction depot use case to illustrate the available options (see Figure 3).

- Physical infrastructure of the shared platform. There are three main options here:
  a. a centralized approach, where the smart contract is only evaluated at one of the banks;
  b. a semi-decentralized approach, where the banks build a shared infrastructure that is conceptually similar to a permissioned blockchain[27];
  c. a fully decentralized approach, where the ST2.0 platform relies on public infrastructure provided by public blockchains (ledgers) such as Bitcoin and Ethereum.

Any smart contract covering a construction depot will necessarily involve a substantial amount of money. It is also likely to involve the disclosure of information about business agreements between the collaborating parties. A fully decentralized platform approach would, therefore, involve serious confidentiality and integrity issues. Given that the builder (or builders) and the buyer may well use different financial service providers, a semi-decentralized infrastructure appears to be a valid option. Alternatively, if a centralized approach is chosen, bilateral agreements will have to be made between the banks involved, and between the banks and the data sources that require interaction with the construction depot's smart contracts. Future research into the physical infrastructure requirements should therefore address the following questions: 1) What are the security requirements involved in implementing a permissioned shared blockchain between Dutch banks? 2) How might confidentiality be assured for smart contracts in a permissioned shared blockchain?

Smart contracts will require trustworthy, secure and transparent interfaces

[27] http://bitscan.com/articles/permissioned-and-unpermissioned-blockchains-part-1

- Smart contract expressiveness and authorship. Smart contracts are Programd in software and may contain costly errors. In a semi-decentralized platform approach, mistakes in smart contracts would be difficult (or, in some cases, nearly impossible) to correct. Given the large amounts involved, and the lengthiness of a smart contract for a construction depot, there is a great need for verifiable, secure and bug-free code. For the same reason, the expressiveness of the smart contract, along with the question of who will be involved in writing it, will become a matter of great importance. This gives rise to the following security issues: 1) What design features would a shared platform need to achieve a balance between expressiveness, security and usability (conveying intent)? 2) When writing smart contracts, how can they be made bug-free and secure? 3) Can any existing software verification technique be used to achieve this goal?
- Interfaces with existing systems. Smart contracts require information from other systems before they can be executed. In the case of construction depots, there will likely be many such information events, including delivery of construction materials and personal, construction-related request and approval messages. These information events will trigger smart contracts to initiate a transfer of value. In this situation, no single party is required to have a complete overview of the system as a whole. Additional controls will therefore be needed to assure trust among stakeholders, to validate new relationships between them and, in general, to prevent malicious actors from exploiting the system. This will most likely require well-specified identity management and PKI. This poses the following problems for security research teams: 1) How can we add accountability to existing information sources to make them secure enough for use in smart contracts? 2) How can we interface with existing value systems to engender sufficient confidence in all parties in the system?

Smart contracts require information from other systems before they can be executed. In the case of construction depots, there will likely be many such information events, including delivery of construction materials and personal, construction-related request and approval messages

User perceptions

With these conditional transaction services looming on the horizon, it is important to understand how users view these concepts and the associated cyber security issues. FinTechs need input of this kind in order to design services that meet users' expectations. With this in mind, we recruited two user groups to help us understand their views and expectations. A preliminary qualitative analysis of the two focus groups, in the context of the use cases presented here, showed that consumers have difficulty with the concept of a 'conditional transaction'. This may be because the respondents assume that end-users do not interface directly with the smart contract. The respondents expect predefined sets of contracts, or at least a template, to cut the time needed to draft their smart contract. We also noted that users were particularly concerned with various aspects of privacy and responsibility. In other words, 'smart contracts' for washing machines should be easy to understand, while personal data and consumption data need to be handled with care. In addition, with regard to the conditions contained in the smart contract, users expect liability to lie with the manufacturer or with the contract's service provider. Finally, respondents want to be able to access all historical and planned transactions in real time. They also want details of the logic used to initiate transactions under the smart contract's provisions.

Conclusion and outlook

Over the past few years, a number of trends have emerged in the transaction web. Some of these are potentially disruptive innovations, such as

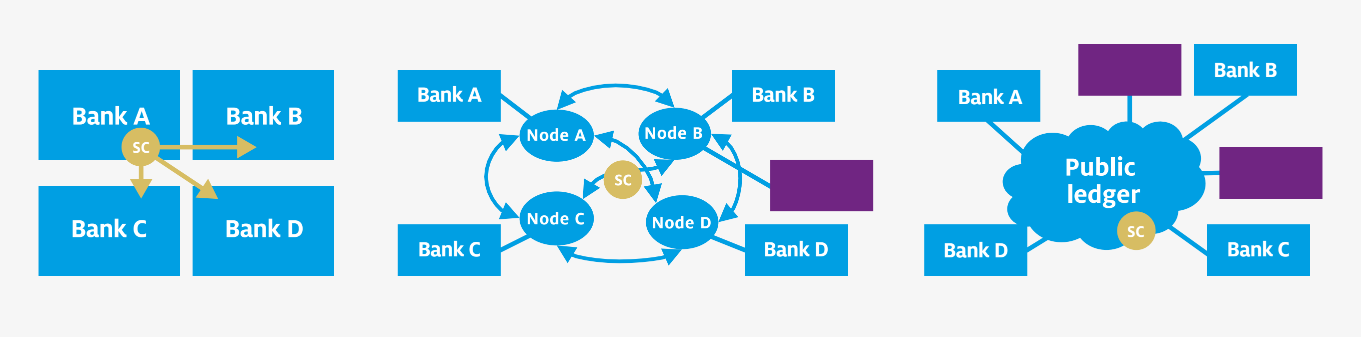

Figure 3: Various options for the distribution of infrastructure in an ST2.0 platform

cryptocurrencies. Others are seen as a natural consequence of developments in today's world, like the unbundling of the transaction chain. If these trends are combined, FinTechs will soon be able to offer revolutionary new transaction services to consumers. These include machine-initiated payments, 'I owe you' groups and automated construction depot services. However these new trends and services will also face new cyber security threats, which need to be identified and carefully addressed. After analysing various potential future services, we identified a number of cyber security issues that require attention. Broadly speaking, these can be divided into two categories: assuring control and maintaining the CIA triad.

The insights we gained in this way have led us to identify a need for conditional transactions – together with a shared ST2.0 service platform to facilitate their realization. The blueprint we developed for this platform allows us to analyse future cyber security challenges for ST2.0.

Future research will initially focus on refining the design and analysis of the shared ST2.0 platform. In particular, we plan to investigate an approach that involves a semi-decentralized infrastructure. This appears to be the most promising design choice, as it strikes an effective balance between retaining a given level of control for individual parties, while reducing the number of bilateral relationships involved in the transaction.

If smart contracts are to be successful, the most pressing issues – such as expressiveness and authorship – will need to be resolved

Innovating in Cyber Security

First phase results of Shared Research Program (SRP)

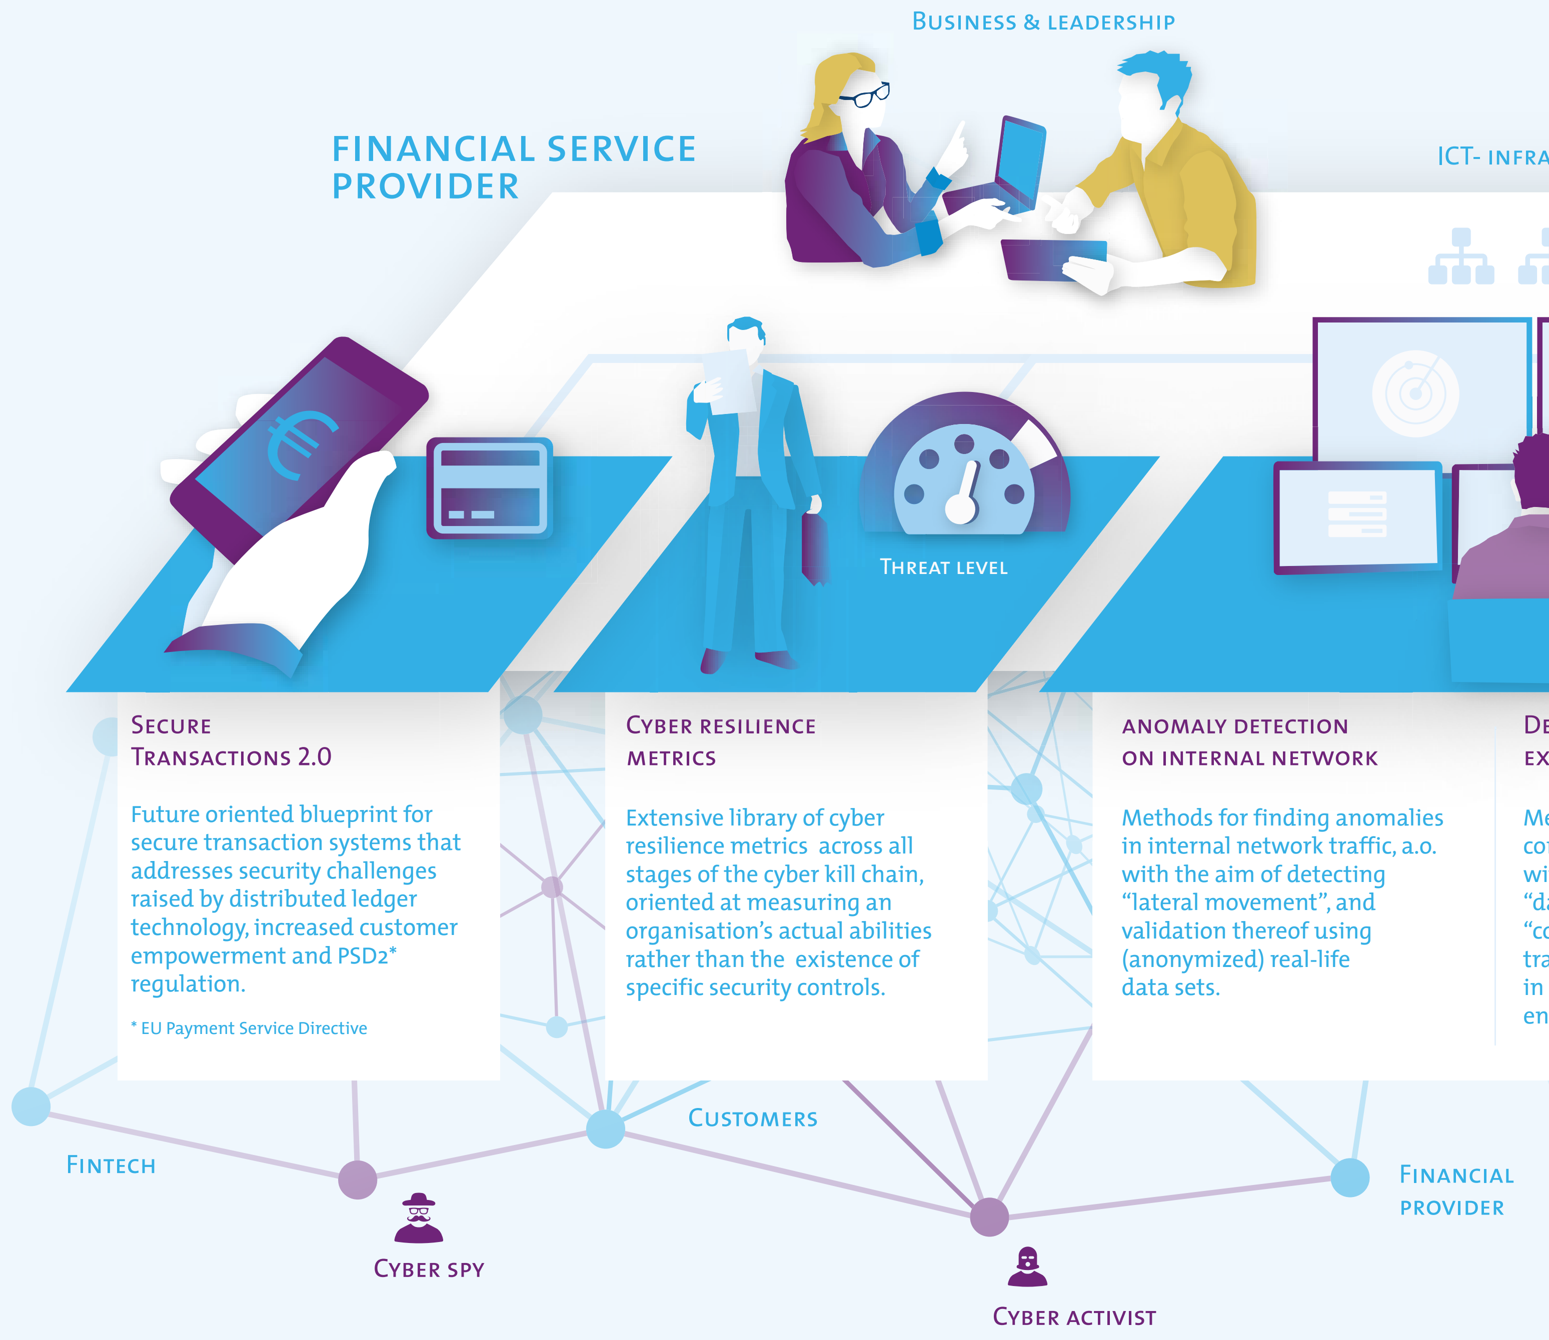

Secure Transactions 2.0

Future oriented blueprint for secure transaction systems that addresses security challenges raised by distributed ledger technology, increased customer empowerment and PSD2* regulation.

* EU Payment Service Directive

Cyber resilience metrics

Extensive library of cyber resilience metrics across all stages of the cyber kill chain, oriented at measuring an organisation's actual abilities rather than the existence of specific security controls.

Anomaly detection on internal network

Methods for finding anomalies in internal network traffic, a.o. with the aim of detecting "lateral movement", and validation thereof using (anonymized) real-life data sets.

De
Ex

Me
co
wi
"da
"co
tra
in
en

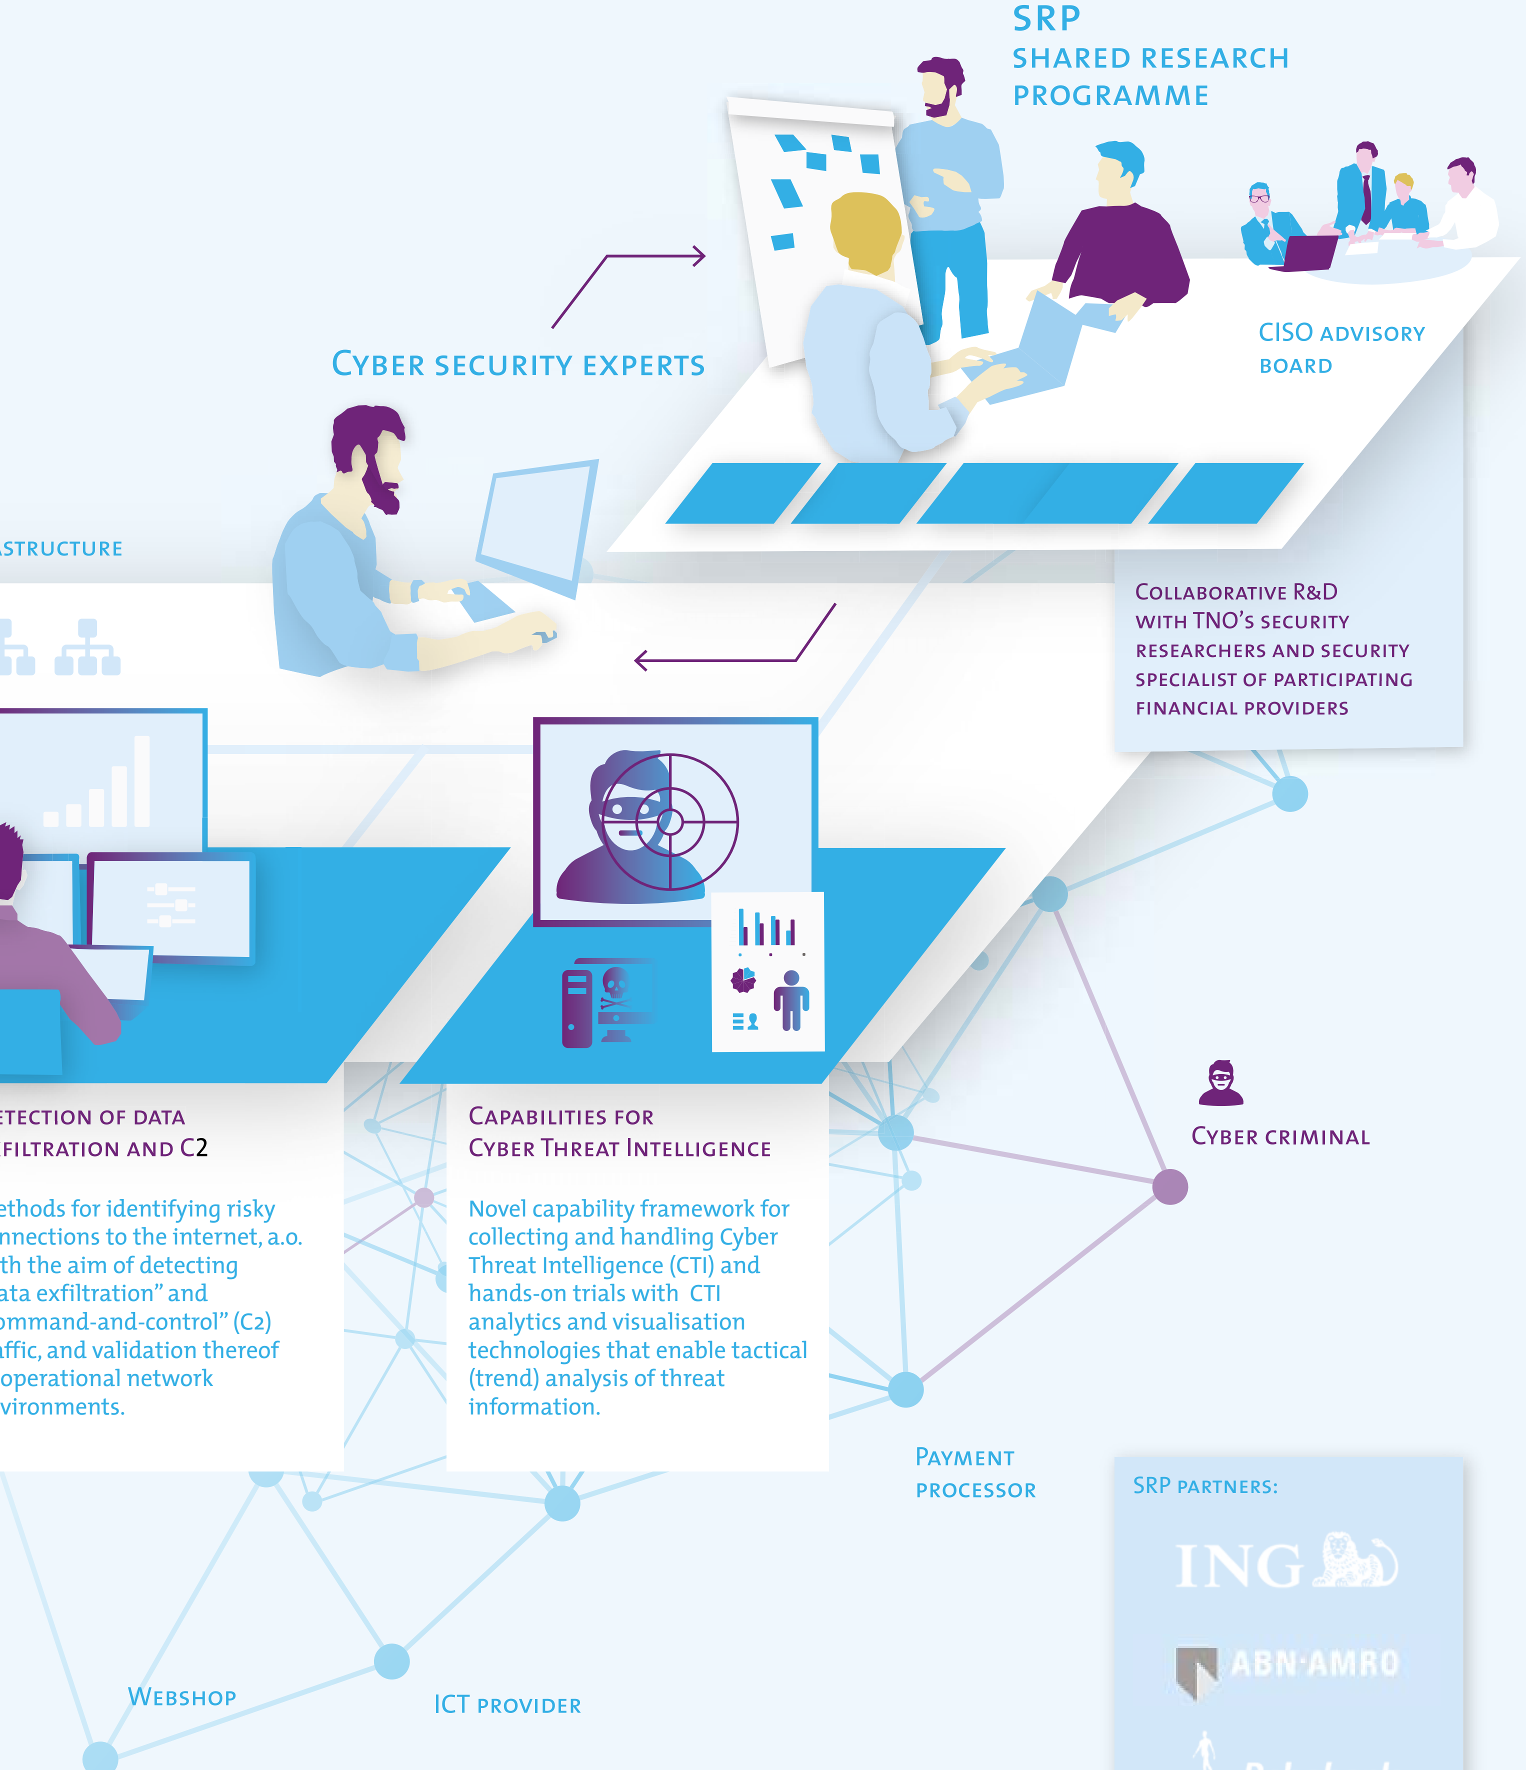

SRP
SHARED RESEARCH
PROGRAMME

Cyber security experts

CISO advisory
board

Collaborative R&D
with TNO’s security
researchers and security
specialist of participating
financial providers

structure

etection of data
filtration and C2

ethods for identifying risky
nnections to the internet, a.o.
th the aim of detecting
ata exfiltration” and
ommand-and-control” (C2)
affic, and validation thereof
operational network
vironments.

Capabilities for
Cyber Threat Intelligence

Novel capability framework for collecting and handling Cyber Threat Intelligence (CTI) and hands-on trials with CTI analytics and visualisation technologies that enable tactical (trend) analysis of threat information.

SRP partners:

ING

ABN·AMRO

Rabobank

TNO

“Collaboration is key in this research field. The Cyber Security Shared Research Program is a fine example of this. TNO is always actively seeking cooperation of this kind. Other examples are the recently launched innovation lab for Cyber Threat Intelligence and TNO’s participation in the European Cyber Security Organization (ECSO).”

Henk-Jan Vink
Director of the Networked Information Roadmap at TNO

Internal network monitoring for targeted attack detection

With billions of euros being transferred online, cybercrime is a very lucrative activity. The recent Swift hack (Riley and Katz 2016) has shown that criminals are capable of executing complex, long-term attacks on financial institutions. However, Dutch financial institutions have now intensified their collaborative efforts in the area of cyber security. This has led to a massive reduction in the damage caused by cybercriminals (NVB, 2016). Internet banking fraud in the Netherlands has actually fallen from a total of €35.1 million in 2011 to €3.7 million in 2015, which illustrates the effectiveness of the financial institutions' cooperative efforts in the cyber security domain. This is forcing cybercriminals to adopt increasingly sophisticated methods of exploiting organizations' vulnerabilities. In this arms race, the only way to stay one step ahead of the criminals is to continuously develop new and innovative methods of preventing and detecting cyber attacks.

Internet banking fraud in the Netherlands has actually fallen from a total of €35.1 million in 2011 to €3.7 million in 2015

Eggshells

The outer layer of defences that separate internal networks from the internet consists of traditional and effective security architectures for mitigating cyber security risks. If a cyber attack is to generate value, it needs to access internal systems containing account balances or other sensitive information, for example. The outer layer of defence is intended to make it more difficult for attackers to penetrate an organization's network. This layer often involves firewalls and IDSs that monitor data traffic. Many security solutions focus on establishing a 'hard' outer layer like this, while leaving the internal network wide open and unmonitored (Shiravi, et al. 2012). This reliance on an outer layer creates what is known as the 'eggshell' security principle (hard on the outside, soft on the inside). Once inside, it is easy for attackers to probe the system at will. So there is a need for methods that detect the malicious activity of attackers who have penetrated the 'hard' outer shell and who now have access to the 'softer' layers within. This can be illustrated by means of a small, fictitious network (see Figure 1) consisting of seven machines, two of which are located in a secure zone. The figure includes various types of malicious internal network traffic based on notorious cyber attacks that have been described in the literature (e.g. Flame, Duqu, Stuxnet).

To complement state-of-the-art techniques, we have developed improved methods for detecting attackers inside the network. One particular type of attack in which internal activity can be observed, is a targeted attack or Advanced Persistent Threat (APT).
Targeted attacks usually:

- involve well-funded, and highly skilled attackers;
- target specific organizations;
- are stealthy and, therefore, difficult to detect.

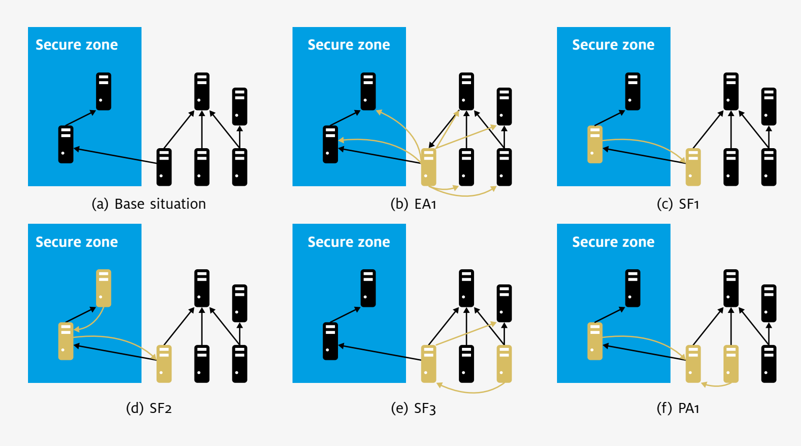

(a) Base situation

(b) EA1

(c) SF1

(d) SF2

(e) SF3

(f) PA1

Figure 1: Examples of malicious internal traffic caused by targeted attacks (Beukema 2016). b) Probing machines for vulnerabilities (Flame, Duqu, Stuxnet), c) Peer-to-peer C&C channel (Duqu), d) Peer-to-peer threat update (Stuxnet), e) Local proxies to steal intelligence (Flame), f) Staging server to gather intelligence (Duqu).

In cooperation with our project partners, we have now developed and implemented a range of new APT detection mechanisms. These are designed to deal with situations in which attackers have already established a foothold inside the organization. During the development phase, we focused on internal network traffic and on pattern-finding in communications between internal hosts. In this context, a host can be any machine that is connected to the internal network. Hosts are identified by their IP addresses (which have been anonymized).

Targeted attacks and internal network traffic

A targeted attack involves a number of steps or phases. The two attack models most commonly reported in the literature are from Lockheed Martin (Hutchins, Cloppert and Amin, 2011) and from Dell SecureWorks (SecureWorks, 2013). Figure 2 illustrates the mapping between these two frameworks, which shows that these two models are quite similar. Both frameworks show that setting up defence mechanisms at the interface between the internal network and the internet can, indeed, prevent or detect many phases of targeted attack campaigns. However, these frameworks also indicate that an important element of a targeted attack campaign takes place within the borders of the targeted network, where it can remain undetected for months or even years. The SecureWorks phases *Expand access* and *Strengthen foothold*, in particular, result in targeted attack behaviour that is rarely detected at the outer border of the network, but which often generates internal network traffic. These phases coincide with the *Command and control* and *Actions on objective* elements of the Lockheed Martin kill chain. Hence internal network traffic is a valuable source of information. By monitoring this traffic, organizations will be better able to detect the later stages of targeted attacks. Our focus, in this project, was to improve detection by monitoring internal network traffic. Current security products do not focus on internal network traffic. Other information sources, such as end-point monitoring, were beyond the scope of this project.

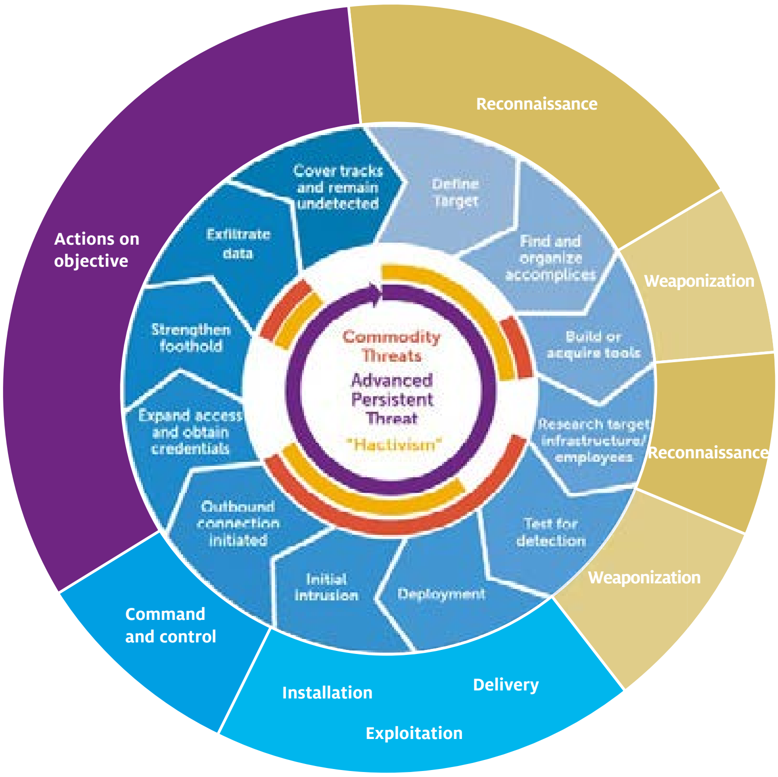

Figure 2: Mapping Dell SecureWorks (inner circle) and Lockheed Martin (outer circle) targeted attack kill chain.

Each model will represent the behavior specific to the associated cluster enabling it to detect all deviations

Monitoring internal network traffic involves specific technical challenges. Network infrastructures typically have a limited number of physical connections to the internet, which makes it relatively easy to monitor the traffic that is passing through these connections. Within a network, however, there could be numerous routes between the various hosts. Thus, monitoring all possible connections can be quite a challenging task. However, the network includes certain locations or concentrators where a lot of information about the hosts' connection behaviour is held. Using these concentrators, we can obtain an overview of the entire internal network by monitoring only a limited number of locations, such as the internal DNS servers or the IAM servers. One remaining key challenge is to extract relevant information without violating any privacy constraints, while keeping the amounts of data involved within manageable limits. Further details concerning the choice of concentrators are given below.

Internal Domain Name Service (DNS) traffic

The first information concentrators to be considered are the internal DNS servers. DNS servers map domain names to IP addresses. Connections are usually preceded by a DNS query that retrieves the IP address. A network typically contains only a few DNS servers (and sometimes only one), which hold details of the DNS queries from each of the hosts. DNS queries do not provide any details about the content of the connection itself, but they do give an accurate picture of how and when hosts are communicating. This means, for example, that – for a given host – we can detect abnormal numbers of interactions with file servers, or suspicious communications with systems that are located in a secure network zone. As we are only interested in internal communication patterns, this approach can be used regardless of whether or not the network has a proxy to perform external DNS queries on behalf of internal hosts.

Identity and Access Management (IAM)

The Identity and Access Management (IAM) service is another location where data concerning hosts' communication behaviour is concentrated. The IAM service processes all requests to access file servers and ensures that only appropriate access is granted to resources within the network. In practice, the high event rate makes IAM logging quite challenging. One possible solution is to limit collection to relevant events. Logging requests to the IAM service shows which file servers were contacted by which hosts and whether the required credentials were used. This information can be used, for example, to detect lateral movement in APTs (Abe 2016).

Internal NetFlow

A network may contain specific central routing nodes that handle significant amounts of traffic. These nodes can be used to monitor that network traffic by means of full packet capture or traffic metadata. Full packet data make it possible to carry out deep packet analysis, which can even be used to inspect details of the network traffic's content. The analysis of full packet captures requires significant processing power and storage capacity. An alternative option is to store and process NetFlow data. NetFlow data is a common way of summarizing network activity. It focuses on meta data and makes it possible to analyse communication behaviour without storing details of the network traffic's content.

Logging requests to the IAM service shows which file servers were contacted by which hosts and whether the required credentials were used

Datasets for behavioural analysis

In our sample network, each host has a host name. Whenever a host wants to communicate with another host, it queries the DNS server for the other host's IP address. No proxy is present. The DNS query logs were acquired by logging the network's DNS server for one week. Table 1 indicates some general statistics for the DNS data, showing that a vast amount of data was generated by the network. The challenge in working with such large volumes of data lies in exposing the relevant information hidden within all that background noise. If the analysis is to be at all useful, preprocessing will be required, to extract the relevant information.

DNS query log	
Date	January 2015
Duration	1 week
Number of Queries	500 million
#Internal A-queries	150 million
#Hosts	>50 000

Table 1: DNS log

Figure 3 shows the total number of internal DNS queries that initiated connections within a one-hour time frame. In this graph the nodes represent the active hosts, and the edges indicate that there has been at least one DNS query from the source node to the destination node. This figure highlights the various tasks carried out by hosts within the network. On the one hand, we can see central hosts that communicate with many other hosts (e.g. file servers) while, on the other hand, there are hosts that communicate with only one or two other hosts. This information enables us to determine the typical behaviour of individual hosts, which in turn makes it possible to detect deviations from this behaviour. A host changing its role within a network might be an indication of malicious activity, such as that seen during the *Expand access* and *Strengthen foothold* phases of a targeted attack.

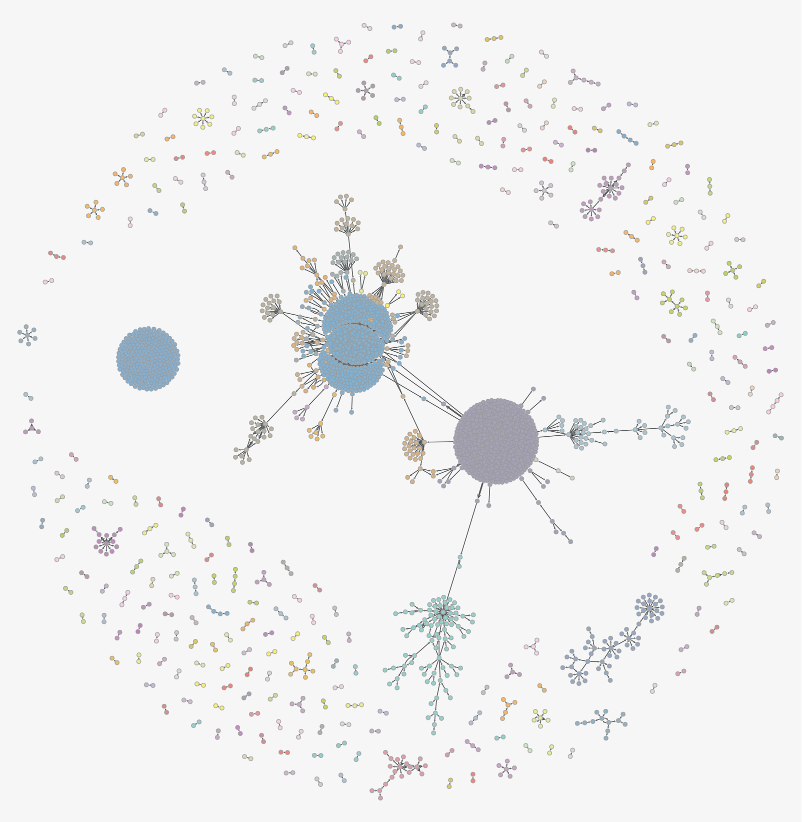

Figure 3: Host activity on 27 January 2015 between 04.00 and 05.00. Within this limited time frame, host activity can be effectively visualized. The nodes represent the active hosts and the edges indicate that there has been at least one DNS query from the source node to the destination node.

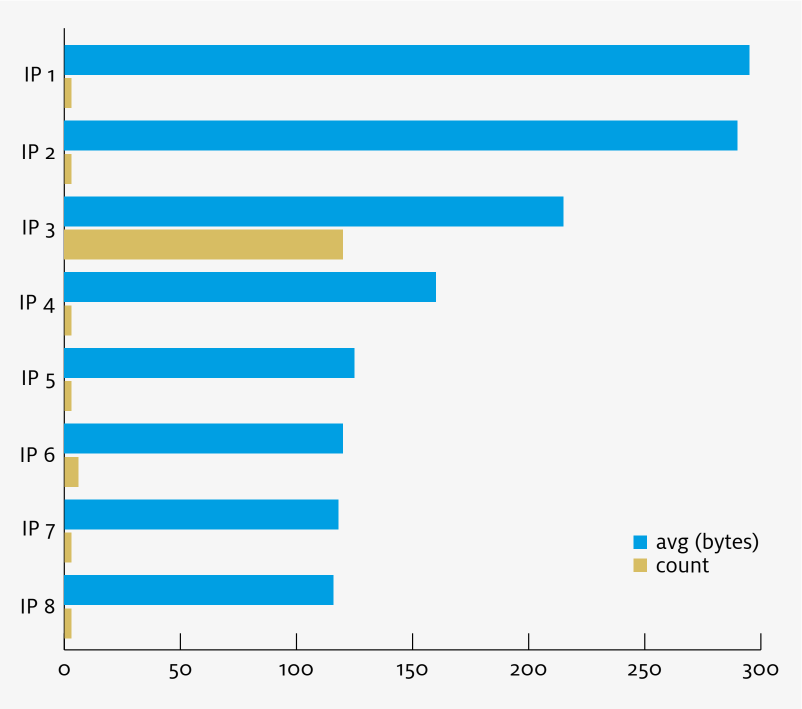

Figure 4: Validation results showing the top nine external IP addresses with the largest average packet size (avg(bytes)), together with the total number of packets (count). The IP 3 address is an experimental DNS tunnel, which is distinct both in terms of the size and number of packets.

In the DNS dataset, both internal hosts and external domains are resolved via the DNS server. In fact, approximately 40% of the A queries in this DNS dataset are the result of an internal host performing a DNS lookup for an external domain and host. So this data source, too, provides interesting information about the communication behaviour of internal hosts with external domains. This inspired us to develop tools to inspect these types of DNS queries, to detect specific techniques that previously have been used in targeted attacks, such as DNS tunnelling and Fast Flux. These detection mechanisms are not based on internal network traffic. Yet, the DNS data generated some directly applicable use cases for detecting those phases in the kill chain that require connections across the network's boundary. DNS tunnels were detected by monitoring external NetFlow data, while Fast Flux detection is based on external DNS traffic. DNS tunnels were found in Carbanak (Carbanak 2015) and, possibly, in Titan Rain. Finally, Fast Flux was seen in Zeus, Flame (Gostev 2012) and in various types of ransomware. For more details see the corresponding text boxes.

Analysing host behaviour

Before any deviations from a host's typical behaviour can be detected, it is first necessary to establish a baseline for the typical behaviour in question. Based on the observed behaviour, a model is created that is capable of estimating the probability that specific events will occur. This model represents a baseline for typical network behaviour. All rare events (which only have a small probability of occurring) may be considered anomalous. These baseline models can be improved by means of clustering techniques.

An approach that involves capturing the behaviour of every host in a single model might prove to be inadequate, as shown in Figure 3. The network contains various types of hosts, which may not all exhibit the same behaviour. Some of these hosts play a pivotal part in the network, initiating many queries. Other hosts are only involved in one or two queries. Experiments have shown that the activity of some individual hosts can be very marginal, which is a problem because a significant amount of data is required to create a statistically sound model. This results in a trade-off between the desire to be able to model hosts as accurately as possible and the need for a sufficiently large

DNS Tunnelling based on external NetFlow

DNS can be exploited to avoid various common security measures. One example is DNS tunnelling, in which data is transmitted via the DNS. Two examples of attacker use cases involving data transfer via DNS are data exfiltration and Command & Control. DNS tunnelling is nothing new, but recently more advanced and stealthy variants have been developed (Ullrich, 2016), (Lee and Schultz, 2016). We have developed an approach to detect DNS tunnelling, based on NetFlow logs. This involves modelling the external DNS servers and identifying anomalies per DNS server, based on the size and quantity of the packets being transferred. We were able to detect an experimental DNS tunnel by analysing the NetFlow data. Using this approach in Splunk, we were able to successfully detect a DNS tunnel in real-time, as shown in Figure 4. Additional details on Splunk implementations can be found in (Jaworski, 2016). Like most DNS tunnelling detection systems, this approach examines the boundary of the network. However, this particular approach can be used at any interface between different sections of a network. It could be used, for example, to check for DNS tunnelling activity between the secure zone (which contains more critical systems or more valuable information) and the rest of the network.

DNS tunnelling is nothing new, but recently more advanced and stealthy variants have been developed

Fast Flux based on external DNS queries

Fast Flux is a technique in which the IP addresses associated with a given domain name change at frequent intervals. In 2016, there was an increase in Fast Flux usage, as some types of ransomware use Fast Flux technology to hide their servers. Many legitimate sites, especially cloud services use multiple IP addresses for a fully qualified domain name (FQDN) as well. But these IP addresses usually fall within the same autonomous system numbers (ASN), or in a small range of ASNs. The IP addresses used in Fast Flux set-ups can originate from all over the world and may have multiple ASNs. We have developed a passive DNS-based approach that builds a network-specific database. The database keeps track of the number of IP addresses and ASNs per domain, and when the domain was last seen in the network data. The algorithm has been successfully validated in an operational environment.

dataset. In an effort to find the optimal balance for such a trade-off, we have introduced a technique for clustering or grouping hosts that display similar behaviour. A behavioural model will then be created for each individual cluster. Host clustering will ensure that each model will be based on a significant amount of data. Moreover, each model will represent the behaviour specific to the associated cluster, enabling it to detect any anomalies.

Clustering techniques

Hosts can be grouped together using mathematical clustering techniques (Jain, 2010). There are many such techniques, each of which groups objects in a different way. Given the dataset involved and the type of anomalies mentioned above, the preferred technique would be graph-based clustering. This technique focuses on the connection behaviour of hosts within a network. It clusters the hosts based on their positions within the graph that represents the network.

One of the most popular algorithms for clustering nodes in a graph is the Louvain method (Blondel, et al. 2008). This technique is effective at detecting communities within a network that are well connected internally but which have only a small number of connections to the nodes of other communities. The popularity of Louvain community detection is based on the availability of efficient implementations and on its applicability to large datasets.

This clustering technique was applied to the network shown in Figure 3. Each identified cluster is assigned a different colour. It is immediately apparent that this technique is very effective at identifying groups of hosts that communicate frequently with one another, but rarely with hosts in other clusters. These rare connections between different clusters stand out as being anomalous. Another graph-based clustering algorithm is Stochastic Block Modelling (SBM) (Mossel, Neeman and Sly, 2012). SBM clusters hosts based on their position within the network. Aside from graph-based clustering algorithms, there is feature-based or vector-based clustering (e.g. k-means (MacQueen, 1967), which is in a class of its own. Here, various features are extracted from the dataset and stored in a vector. Features that might be of interest in this context are the amount of traffic that is sent or received by a host, the duration of its connections, or the types of DNS queries involved.

Each model will represent the behaviour specific to the associated cluster, enabling it to detect any anomalies

Clustering and anomaly detection

When the desired clustering technique has been applied, the behaviour of nodes within a cluster can be modelled. Any deviations from this behaviour will be detected and considered anomalous. Figure 5 illustrates the general approach we can apply to various data sources. Each data source requires its own specific preprocessing module, which supplies input to the more general clustering modules. When examining the DNS logs, we used every individual step in this framework. We constructed an anomaly detector capable of detecting any hosts that display behaviour which deviates from normal cluster behaviour (based on internal DNS queries). We have observed that the types of anomalies detected are heavily dependent on the type of clustering algorithm selected. Our next step will be to validate the anomaly detector we have developed. Here, anomalies detected will be analysed in cooperation with systems administrators who have an in-depth knowledge of the network in question. Meanwhile, the approach has now been developed to the point at which it can be applied to the other data sources suggested.

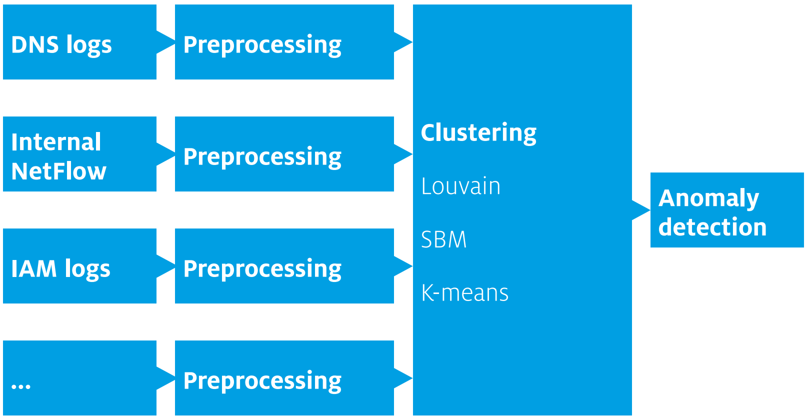

Figure 5: Integrated cluster-based anomaly detection approach. These frameworks show that – following data preprocessing – a range of data sources can be used as input for these clustering techniques. The types of anomalies detected are heavily dependent on the type of clustering selected.

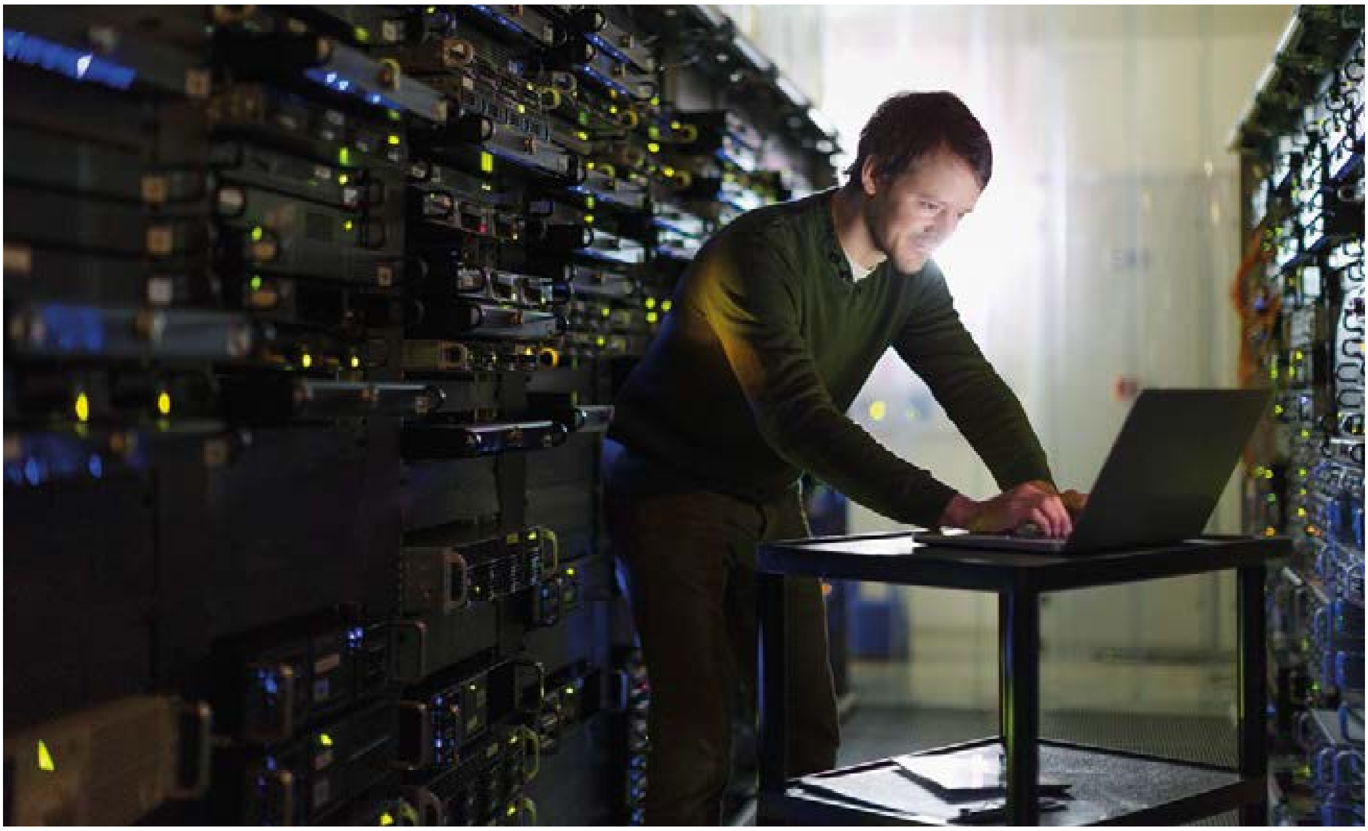

Classification algorithm

An alternative approach to the one shown in Figure 5 is a client-server classifier. This technique enables us to determine the apparent role (client or server) of a host within a network, based on its behaviour. When combined with details of the host's actual role (if that is known), this classifier can be used to determine whether or not the host is behaving as expected. For example, servers are mainly expected to connect to internal hosts as opposed to clients, which also initiate many connections to hosts outside the network. If this technique classifies a client as a server, or vice versa, the machine in question needs to be inspected.

In this research project, we developed a client-server classifier that is based on DNS log information. It extracts a range of statistical features in the preprocessing step. Next, a Kolmogorov-Smirnov test was used to systematically identify the features that distinguish clients from servers. One feature that can be used to effectively distinguish clients from servers is the frequency with which a host is looked up in the DNS resolver. Another such feature is the number of PTR queries per host. A total of six largely independent features were selected, for optimal performance. These features were then used to train a Support Vector Machine (SVM), which could then use new data about a host to classify the host's behaviour as client-like or server-like. The algorithm's performance, in terms of anomaly detection, was analysed by checking whether the classification it generated matched the host's official label (server or client).

Table 2 shows the performance of the client-server classification algorithm. It correctly classified more than 96% of the clients and more than 99% of the servers. The algorithm also detected a small number of anomalous (misclassified) hosts that were not behaving as expected. In addition, the algorithm generates a classification or anomaly score indicating the extent to which a host is deviating from its expected behaviour. This anomaly score allows us to prioritize the anomalies in question, and even to balance the ratio of true and false positives.

Servers are mainly expected to connect to internal hosts as opposed to clients, which also initiate many connections to hosts outside the network

SVM classification	Classified as clients	Classified as servers
Clients	99.6 %	0.4 %
Servers	3.2 %	96.8 %

Table 2: Performance of SVM classification system trained on three consecutive days and evaluated on day 4

Clustering techniques proved effective in detecting anomalous internal network traffic. The new cluster-based detector is now ready for every-day use in the detection of Advanced Persistent Threats

Conclusion

One major security issue concerns attackers who, after penetrating the hard security perimeter, establish a foothold inside the organization's network. How can they be detected and defeated? Traditional security solutions focus on detecting and preventing the initial steps of an attack. As a result, if an attacker succeeds in breaking through this defence, the internal network is entirely unprotected.

In this project, we took a different approach, focusing instead on data sources that spotlight a network's internal traffic. We developed a novel and generic anomaly detection framework that was based on clustering techniques. To test this approach, we applied it to internal DNS traffic. Clustering techniques proved effective in detecting anomalous internal network traffic. The new cluster-based detector is now ready for everyday use in the detection of Advanced Persistent Threats. Tests involving internal DNS traffic showed that DNS servers are useful concentrators for detecting anomalous network traffic.

We also developed a client server classification algorithm (which deviates slightly from the framework) that is now ready for further validation in operational environments. We also took some of the by-products from this project and developed them into readily applicable detection techniques for use in the later stages of an attack. One involves data exfiltration, using DNS, while the other is based on command and control traffic, using Fast Flux techniques. The DNS tunnel detector and the Fast Flux detectors make use of other data sources (external NetFlow and external DNS traffic, respectively), and involve information about different aspects of a host's communication behaviour.

“The threats facing the banking industry are changing rapidly. Our security measures are also constantly evolving, with a view to defending our customers. This evolutionary process is driven by constant innovation in the areas of cyber crime prevention, detection and response. One aspect of this innovation is our participation in the Shared Research Program. We feel that this combined effort will improve security for our customers and promotes joint innovation across a broad range of countermeasures against cyber crime.”

Vincent Thiele
Manager CCERT ING

Towards a mature cyber threat intelligence practice

The cyber threat landscape has undergone an enormous evolutionary leap. High-end cyber attacks are now carried out by professional organizations with advanced technical capabilities and substantial resources at their disposal. Such attacks are often targeted and persistent and they can involve great technical sophistication. In response to the nature and dynamics of present day cyber threats, many organizations have fundamentally revised their cyber resilience strategies. Most prominently, it has become common to complement traditional, preventive security controls such as access control and data encryption with elaborate measures for security monitoring[10] and incident response. This development was instigated by the widely held view that preventive measures cannot avert a security incident if the adversary is sufficiently motivated and competent.

While advancements in monitoring and response have greatly helped to reduce the damage done by cyber attacks, such reactive strategies are not always ideal. In an effort to regain the initiative, many organizations are now developing Cyber Threat Intelligence (CTI) capabilities. In essence, the idea is to anticipate emerging cyber threats rather than await an actual incident. Collecting and handling CTI is a relatively new field, so many of the practices and solutions involved are still at the pioneering stage. In light of this, the SRP partners explored what it would take to establish an effective and mature Cyber Threat Intelligence practice.

This article introduces the concept of CTI and explains some of the key insights generated by the program. Most importantly, it presents the framework for CTI capabilities that resulted from the first phase project, the needs that this framework intends to fulfil and the lessons learned from the joint design effort.

The CTI playing field

There is no commonly agreed definition of the term ‘Cyber Threat Intelligence’. However, the description formulated by Gartner[11] covers the basic essentials:

> “Threat intelligence is evidence-based knowledge, including context, mechanisms, indicators, implications and actionable advice about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject’s response to that menace or hazard.”

Actual CTI can take many forms. Typical examples include:

- Indicators of Compromise (IoCs). This refers to noticeable artefacts on hosts or networks that are indicative of malicious tools or known methods of attack (TTP, see below). These may be evidence of an ongoing intrusion.
- Tactics, Techniques and Procedures (TTPs). TTPs reflect the Modus Operandi (MO) of cyber adversaries. This includes their actions at specific stages of an attack, the tools and techniques employed, the resources (infrastructure, personas) that they leverage in the target environment, etc.
- Threat actor profiles. These are the characteristics of cyber adversaries, e.g. identity or alias, objectives/ motivators, typical TTP/MO, known (historical) attacks, suspected associations with other threat actors, etc.
- Attacker campaigns. The characteristics of related attacks and intrusions through which an adversary pursues a larger, strategic goal. These may include attributes of the threat actor (or actors) involved, their suspected objectives, the TTPs employed and details of any related incidents.

[10] The SRP also examined techniques for detecting advanced cyber attacks. This is covered in the article ‘Internal network monitoring for targeted attack detection’ included in this publication.

[11] Gartner, Definition: Threat Intelligence, 16 May 2013, https://www.gartner.com/doc/2487216/definition-threat-intelligence

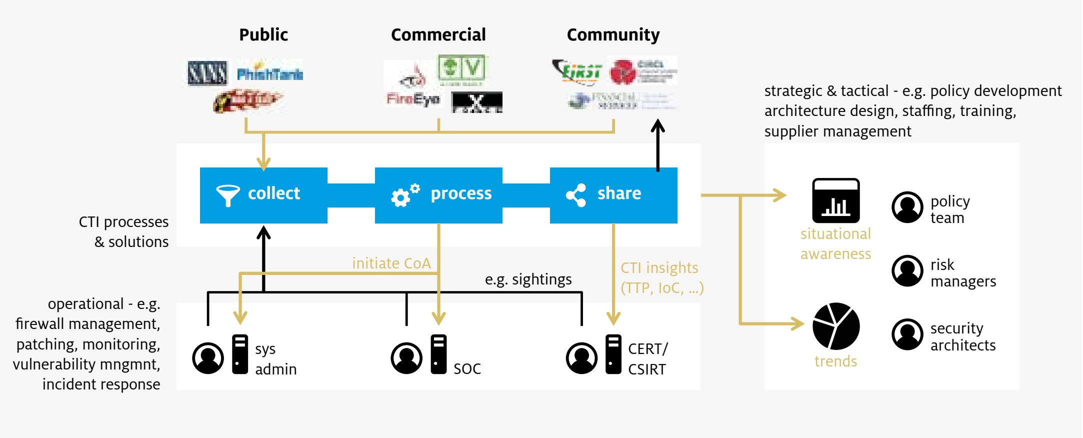
Public
Commercial
Community
strategic & tactical - e.g. policy development architecture design, staffing, training, supplier management
collect
process
share
CTI processes & solutions
initiate CoA
e.g. sightings
CTI insights (TTP, IoC, ...)
situational awareness
policy team
risk managers
security architects
trends
operational - e.g. firewall management, patching, monitoring, vulnerability mngmnt, incident response
sys admin
SOC
CERT/ CSIRT

Figure 1: The CTI playing field

Traditional CSIRT service descriptions do not fully capture the field of CTI

True CTI includes appropriate context and allows an organization to take decisions on the follow-up required. Figure 1 shows how CTI measures are typically positioned in an organizational context.

CTI operations are coordinated by one or more CTI analysts. At present, such analysts are usually members of the organization's CERT[12] or CSIRT[13] team[14]. One of their key duties is to collect internal and external intelligence that might reveal potential threats to the organization. Such threat information can be obtained from a great variety of (public, commercial or community) sources. This typically generates vast amounts of data, some of which is structured (i.e. in a standardised, machine readable format), while some can also be unstructured (narrative, e.g. e-mail advisories or threat reports). The CTI analyst processes this data from two distinct perspectives:

- At the operational level, analysts scan the continuous stream of threat information for the presence of imminent or emerging threats to the organization. If they discover a potentially relevant threat, the analysts initiate an appropriate Course of Action (CoA). This usually means that actions are fed into the organization's operational security processes.
- At the tactical level, analysts compile aggregated statistics and trends in relation to the body of threat information that was collected over time. Such tactical data is fed into mid-term and long-term (i.e. strategic and tactical) security planning processes.

A CoA at the operational level usually takes the form of a real-time response, e.g. a modification of firewall rules or the on-boarding of threat indicators in detection (SIEM[15], IDS[16]) solutions.

Tactical follow-up, on the other hand, will often require careful planning and budgeting. It may, for instance, involve changes in the capabilities of security teams (requiring a reappraisal of recruitment and education strategies) or in the organization's network designs and security architectures (possibly resulting in major migration projects).

Finding a proper footing

The financial providers participating in the SRP all have some level of capability in place for collecting and handling Cyber Threat Intelligence (CTI). While some have made more progress than others, all are still in the early stages of establishing an effective and sustainable CTI practice. Notably, all are actively pursuing enhancement of their CTI capabilities. For instance, all SRP partners are in the process of adopting a dedicated CTI platform to facilitate automation and analytics in their CTI operations. On the whole, however, the present status of CTI activities among the SRP participants reflects the pioneering nature of this relatively new field (as mentioned above).

This raises the question what exactly constitutes a truly 'mature' CTI practice. The SRP partners found that traditional CSIRT service descriptions do not fully capture the field of CTI. Services of this kind were first described in the CERT/CC Handbook for Computer Security Incident Response Teams. This framework has been adopted by various organizations, including the European Union Agency for Network and Information Security (ENISA).
It includes a specification of 'proactive services' (see Figure 2). However, the handbook does not address the concept of Cyber Threat Intelligence, nor does it include particular services that appeal to the CTI playing field described in the previous section.

[12] Computer Emergency Response Team
[13] Computer Security Incident Response Team
[14] In reality, the collection and handling of CTI is often a secondary task assigned to teams whose core priority is incident response. This is not an ideal foundation for a sustainable CTI practice. As the field matures, it is conceivable that dedicated CTI teams will be appointed.
[15] Security Incident and Event Management
[16] Intrusion Detection System

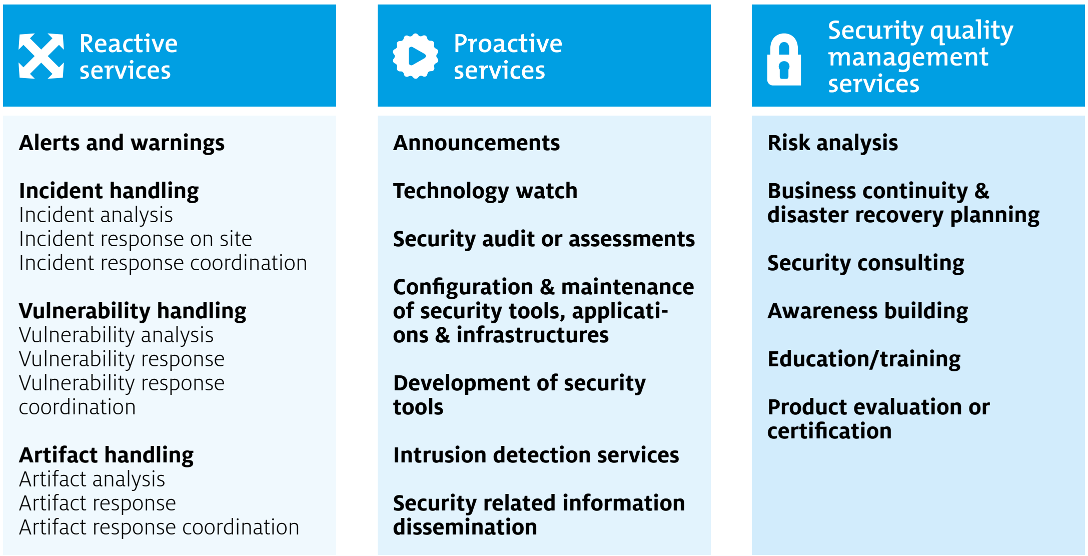

Reactive services

Alerts and warnings

Incident handling
Incident analysis
Incident response on site
Incident response coordination

Vulnerability handling
Vulnerability analysis
Vulnerability response
Vulnerability response coordination

Artifact handling
Artifact analysis
Artifact response
Artifact response coordination

Proactive services

Announcements

Technology watch

Security audit or assessments

Configuration & maintenance of security tools, applications & infrastructures

Development of security tools

Intrusion detection services

Security related information dissemination

Security quality management services

Risk analysis

Business continuity & disaster recovery planning

Security consulting

Awareness building

Education/training

Product evaluation or certification

Figure 2: CSIRT services specified by CERT/CC

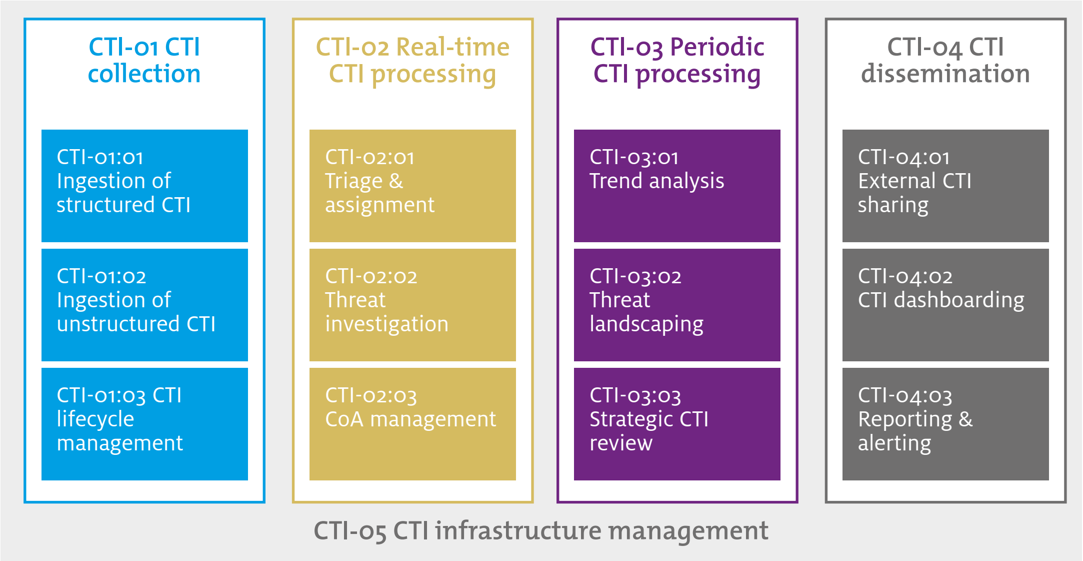

CTI-01 CTI collection
CTI-01:01 Ingestion of structured CTI
CTI-01:02 Ingestion of unstructured CTI
CTI-01:03 CTI lifecycle management

CTI-02 Real-time CTI processing
CTI-02:01 Triage & assignment
CTI-02:02 Threat investigation
CTI-02:03 CoA management

CTI-03 Periodic CTI processing
CTI-03:01 Trend analysis
CTI-03:02 Threat landscaping
CTI-03:03 Strategic CTI review

CTI-04 CTI dissemination
CTI-04:01 External CTI sharing
CTI-04:02 CTI dashboarding
CTI-04:03 Reporting & alerting

CTI-05 CTI infrastructure management

Figure 3: CTI capability framework

MITRE's Ten Strategies of a World-Class Cybersecurity Operations Centre (see literature) offers a more contemporary view of cyber security capabilities and includes a CTI-oriented 'Intel and Trending' element. However, the capability descriptions contained in this document are of a fairly high abstraction level. What's more, MITRE positions CTI as one of several fields in a broadly oriented Security Operations Centre (SOC). Accordingly, generic capabilities – such as reporting – have not been considered specifically from a CTI viewpoint.

Since existing literature did not offer the desired foundation, the SRP partners set out to design a novel CTI Capability Framework, using MITRE's Intel and Trending capabilities as a source of inspiration. The underlying idea was that such a framework would provide a shared understanding of the measures required, while also allowing partners to refine their individual plans for developing a mature CTI practice.

A framework for CTI capabilities

Figure 3 shows the top-level structure of the framework developed by the SRP participants. In essence, it consists of twelve CTI capabilities across four distinct categories. In accordance with the CTI environment outlined above, some of these capabilities are operational in nature (categories CTI-01 and CTI-02) whereas others (CTI-03 and CTI-04) serve a strategic or tactical purpose. Note that CTI Infrastructure Management[17] (CTI-05) was not described in great detail, as its setup will likely be tailored to the specific IT processes of individual organizations.

[17] This capability encompasses activities such as acquisition, maintenance and release management for software solutions and other technical infrastructure employed in the organization's CTI practice.

The capabilities outlined in the framework were broken down into their individual core operations, the envisaged workflows and the information sources (both internal and external) that would need to supply input. A critical design choice in this respect was to fully detach the framework from common security functions and security team structures. The duties and organizational placement of SOC facilities and CSIRT teams, for instance, tend to vary from one organization to another. Thus such terms do not bear the exact same meaning in all situations. In light of this, they were deliberately avoided when designing the CTI capability framework.

Interestingly, some elements of the framework were more or less anticipated in advance, whereas others were less obvious. This is illustrated by the first category of capabilities, which involves the collection of appropriate threat information (CTI-01, see Figure 3). Here, CTI-01:01 covers the ingestion of CTI that is available in standardized (machine-readable) formats[18]. This capability is defined as follows:

> The ability to consume, normalize and enrich machine-readable threat information and feed it to the organization's centralized CTI repository in a fully automated fashion.

The SRP participants were already aware of the need for a capability of this kind. In fact, all were already ingesting automated CTI feeds from a range of (mostly commercial) sources. What's more, centralised CTI repositories – while not common just yet – should naturally result from the already initiated adoption of dedicated CTI platforms (see above). An insight that was less obvious upfront, however, is that organizations can also ingest structured CTI from internal sources. The malware analysis tools used in incident investigations can for instance generate threat indicators. Similarly, the usually present SIEM solutions can reports so called 'sightings' (i.e. actual hits on an IoC). Information of this kind is a valuable addition to the organization's CTI repository, but active collection of such data is not exactly common practice. Designing this framework therefore gave the partners a better understanding of what this capability might entail.

CTI analysts are faced with ever increasing volumes of threat information and the need to respond to threats ever more rapidly

The same was true of CTI-01:02, which covers the ingestion of unstructured CTI. This capability is defined as follows:

> The ability to deduce machine-readable characteristics from narrative or otherwise unstructured threat information and feed these to the organization's centralized CTI repository.

Unstructured CTI is threat information that is not shaped in a standardized markup format. Rather it takes the form of a .pdf documents or the body text of an e-mail. Typical examples include expert reports and advisories of the kind produced by government agencies, solution vendors and specialized investigative companies. The organizations involved in the SRP were already ingesting threat information of this kind. Converting such threat information to machine-readable records in a centralized repository, however, was not common practice. Interestingly, the value of such conversion became apparent during technical trials that involved capabilities of a more tactical nature. This will be explained in the following section.

Another less obvious capability was CTI Life-Cycle Management (CTI-01:03), one aspect of which is the periodic review of threat information sources:

> The ability to ensure that the sources of CTI employed by the organization adequately meet the information needs of CTI analysts, and those of any stakeholder involved in CTI handling, at any given time.

Many organisations want to know which are the 'best' sources (either public or commercial) of threat information. In reality, however, there is no universally applicable silver bullet that suits the context of each individual organization. The SRP partners gradually came to realize that maintaining an appropriate set of CTI sources is, in fact, a capability in itself. Such life-cycle management includes periodic reviews of the CTI sources that the organization has been ingesting. This is best achieved by routinely collecting statistics that reflect the quality and relevance of individual

[18] Typical examples include STIX documents and .csv files.

threat information feeds. Relevant metrics in this regard include the extent to which a source supplied unique threat insights (not reported by any other source) or enabled the discovery of actual attacks or intrusions (rather than producing large numbers of false positives). Ideally, the process of compiling such performance statistics would be largely automated. As yet, however, few organizations have acknowledged the need to establish a capability of this kind[19].

All in all, the process of jointly developing a capability framework gave the partners a better understanding of the practices they would need to put in place to exploit the field of CTI to its full potential.

Experimenting with CTI analytics

Automation is widely acknowledged to be a key facilitator in the optimization of CTI operations. CTI analysts are faced with ever increasing volumes of threat information and the need to respond to threats ever more rapidly. What's more, the process of acquiring valuable insights from such threat information requires increasingly complex analytics that transcend the capabilities of human experts. The SRP partners explored a range of tools and technologies that were potentially capable of supporting the trend analysis capability (CTI-03:01, see Figure 3). Focusing on a tactical capability of this kind was considered most valuable since the financials believed that automation needs at the operational level would be resolved by the adoption of a dedicated CTI platform (which all had already set into motion, see above).

To acquire some actual hands-on experience, the project established a Proof of Concept (PoC) environment for analysing CTI statistics. The core of this environment consisted of an open source tool chain[20] for data analysis and visualization. This tool chain was fed with a substantial volume of threat information (mostly indicators), obtained from various public sources. This setup was subsequently used to analyse a variety of contemporary threats, some general in nature, others derived from internal investigations carried out by individual SRP participants. Figure 4 shows an example of the statistics and visualizations that were compiled. This particular graph traces the evolution of a family of ransomware[21] in the first few months of 2016.

Analysing threat information in TNO's Cyber Threat Intelligence Lab

The experiments revealed that statistical analysis of threat information is feasible and offers genuine value. Although acquiring entirely new insights (e.g. concerning the severity and evolution of threats) proved difficult, the statistics and

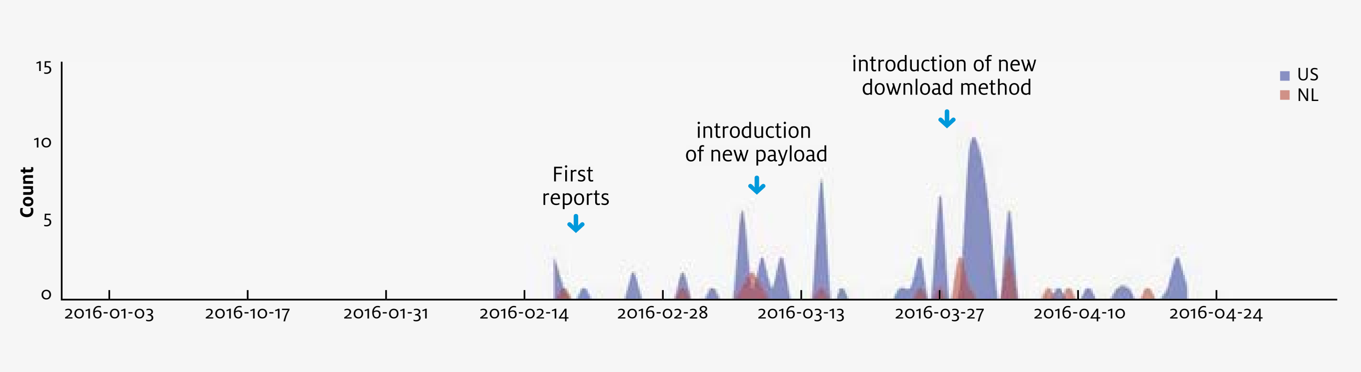

Figure 4: Evolution of 'Locky' ransomware [US activity depicted in blue, NL activity in red]

[19] One of the SRP participants actually evaluated the performance of its CTI sources well before this capability was included in the framework. While this evaluation was ad hoc and manual, it proved to be an effective foundation for rationalizing ('cleaning up') the overall amount of CTI ingested.

[20] To this end the project used the popular Elastic Stack, see https://www.elastic.co/products.

[21] Computer malware that holds a victim's data hostage or threatens to publish it until a ransom is paid.

visualisations did substantiate expert opinions (or suspicions) with tangible evidence. The SRP partners believe this can provide a stronger basis for formulating defensive strategies and justifying the corresponding investments.

On the whole, statistical analysis of public threat information supplied some surprisingly useful insights. Incorporating CTI compiled by an actual financial provider should enhance the results even further. However, threat data of SRP partners proved to be impractical for experimental purposes, as this information is presently scattered across a range of information systems and is largely unstructured in nature. Thus the operational setup that was initially believed to be satisfactory (see above) actually limited the development of tactical capabilities. This was, in fact, one of the major lessons learned (see the following section) in the course of the project.

Lessons learned

Threat Intelligence has a long history in the military field, both in theory and in actual practice. Within the cyber realm, however, it is a relatively new domain that has not yet been subject to the same level of systematic investigation. The most important lessons drawn from this joint project were as follows:

- The maturity of cyber threat intelligence practices can and must improve. At present, activities in this field are largely operational in nature. More often than not, the approach is limited to the collection of malware samples and Indicators of Compromise (IoCs) and a relatively unstructured dissemination of such threat information to (security and IT) operations teams. Arguably, many current activities in this field are more about sharing information than about conducting actual intelligence work.
- The project showed that improvement of strategic and tactical CTI capabilities requires a solid basis of operational threat intelligence measures. In most cases, however, this operational foundation was not as strong as had been initially thought. Indeed, the SRP partners found that all levels of CTI handling have room for improvement. One particularly telling operational limitation is that the most valuable (tactical) intelligence is often resides in the mailboxes of individual specialists. This effectively rules out any possibility of conducting automated aggregations or analytics.
- Joint sessions involving the banks and TNO revealed gaps in CTI and security operations that were not obvious upfront. An example stems from the newly identified need to cross reference threat information with incident data. This proved to be something of a challenge, as incident data is not always stored in a single location. While the scattered storage of incident data was already considered to be less than ideal, it was not perceived as a material issue up until that moment.
- The capability framework provides a solid foundation on which to build a mature Cyber Threat Intelligence practice. The actual approaches used to establish the various capabilities in practice will likely vary from one organization to another. It might be useful to complement the model with an audit framework and, possibly, some indication of maturity levels, as an aid to actual implementation. In other words, the CTI capability framework has potential for further enhancement.

Arguably, many current activities in this field are more about sharing information than about conducting actual intelligence work

The SRP proved to be a useful platform for discussing the ‘how’ of Cyber Threat Intelligence (above and beyond sharing operational content such as IoCs). It was this cooperation that enabled the financials to leverage their joint expertise and capabilities in the area of CTI. Despite the obvious importance of this field (see Introduction), few organizations actually have a mature CTI practice in place. Now is a good time for organizations to determine their exact status, and to pursue enhancements. The SRP partners will certainly be doing so. Given that the partners are all at a similar level, there is great potential for further collaboration in this area.

“TNO is ideally suited to the Cyber Security Shared Research Program because its core competence is applied research.
It takes results from scientific research at universities and translates them into usable concepts and pilot products for end users, such as financial institutions. This is exactly what the Cyber Security SRP is all about.”

Annemarie Zielstra
Director of Cyber Security and Resilience at TNO

Literature

Advanced persistent threat analysis, Dell SecureWorks, 2013. Available: https://www.secureworks.com/blog/advanced-persistent-threats-apt-a

A step-by-step approach on how to set up a CSIRT- Including examples and a checklist in form of a project plan, Deliverable WP2006/5.1 (CERT-D1/D2), ENISA, 2006. Available: https://www.enisa.europa.eu/publications/csirt-setting-up-guide

A. Gostev, The Flame: Questions and Answers (blog), Kaspersky Lab, 2012. Available: https://securelist.com/blog/incidents/34344/the-flame-questions-and-answers-51/ (accessed October 7, 2016).

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

A. Shiravi et al., Toward developing a systematic approach to generate benchmark datasets for intrusion detection, Computers & Security, 2012. Available: http://www.sciencedirect.com/science/article/pii/S0167404811001672

A.K. Jain, Data clustering: 50 years beyond K-means, Pattern Recognition Letters, 2010

C. Zimmerman, Ten Strategies of a World-Class Cybersecurity Operations Center, MITRE, 2014. Available: https://www.mitre.org/sites/default/files/publications/pr-13-1028-mitre-10-strategies-cyberops-center.pdf.

Carbanak APT - The Great Bank Robbery, Kaspersky Lab, 2015. Available: https://securelist.com/files/2015/02/Carbanak_APT_eng.pdf

E. Mossel, J. Neeman and A. Sly, Reconstruction and estimation in the planted partition model, UC Berkeley, Springer, 2014. Available: http://link.springer.com/article/10.1007/s00440-014-0576-6

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

Factsheet Veiligheid en fraude, Nederlandse Vereniging van Banken, 2016. Available: https://www.nvb.nl/media/document/001487_nvb-factsheet-veiligheid-en-fraude-2017.pdf

J. Ullrich, Command and Control Channels Using AAAA DNS Records (blog), Internet Storm Center, 2016. Available: https://isc.sans.edu/diary/Command+and+Control+Channels+Using+%22AAAA%22+DNS+Records/21301

J.B. MacQueen, Some Methods for classification and Analysis of Multivariate Observations, Proceedings 5th Berkeley Symposium on Mathematical Statistics and Probability, University of California Press, 1967

M. Eeten, Economics of cybersecurity - Measuring security levels, Delft University of Technology (MOOC course material), 2015.

M. J. West-Brown et al., Handbook for Computer Security Incident Response Teams (2nd edition), Carnegie Mellon University, 2003. Available: http://www.cert.org/archive/pdf/csirt-handbook.pdf.

M. Lee and J. Schultz, Detecting DNS Data Exfiltration (blog), Talos, 2016. Available: http://blog.talosintel.com/2016/06/detecting-dns-data-exfiltration.html (accessed 08 29, 2016).

M. Riley and A Kat, Swift Hack Probe Expands to Up to a Dozen Banks Beyond Bangladesh, Bloomberg Technology, 2016. Available: http://www.bloomberg.com/news/articles/2016-05-26/swift-hack-probe-expands-to-up-to-dozen-banks-beyond-bangladesh

S. Abe, Detecting Lateral Movement in APTs - Analysis Approach on Wndows Event Logs, JPCERT/CC - ICS security Response Group, 2016. Available: https://www.first.org/resources/papers/.../FIRST-2016-105.pdf

S. Jaworski, Using Splunk to Detect DNS Tunneling, SANS Institute, 2016. Available: https://www.sans.org/reading-room/whitepapers/malicious/splunk-detect-dns-tunneling-37022

V.D. Blondel et al., Fast unfolding of communities in large networks, Journal of Statistical Mechanics, 2008

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

W.J.B. Beukema, Enhancing Network Intrusion, Master thesis, University of Twente, 2016.

Colophon

© 2017 Participants in the Cyber Security Shared Research Program.

Editors

Rachelle Harel, Richard Kerkdijk and Reinder Wolthuis

Contributing experts

ABN AMRO: Bert Heijnen, Roger Lagarde, Nyncke van de Linde, Ilona Scheerder, Mark Wiggerman

Achmea: Tom Huitema, Kadir Kalayci

ING: Ina Eising, Marleen van Emmerik, Jacco Ligthart, Wim Tip, Dave Woutersen, Sander Zeijlemaker, Ruud Zwan.

Rabobank: Jethro Cornelissen, Kees Hoogewoning, Eric Kaasenbrood, Kelvin Rorive, Vincent Spoel.

TNO: Maarten Everts, Frank Fransen, Rachelle Harel, Ruben ten Hove, Jan Kazemier, Sander de Kievit, Joke Kort, Erik Meeuwissen, Erwin Middelesch, Bert-Jan te Paske, Harm Schotanus, Sven Schultz, Eldine Verweij.

Copyright

Contact

reinder.wolthuis@tno.nl
wegwijzer@tno.nl

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