Ontology-based Cybersecurity and Resilience Framework

Helmar Hutschenreuter, Salva Daneshgadeh Çakmakçı, Christian Maeder and Thomas Kemmerich
University of Bremen, Germany

Keywords: Resilience Framework, Security Ontology, Maritime Port Cybersecurity, Ontology-based Response and Recovery, SIEM, Inference System.

Abstract: In the digital age, almost all organizations have become dependent on Information Technology (IT) systems at different levels of their individual and collective activities. Physical infrastructures are inextricably tied to the functioning of IT systems that are vulnerable to internal and external cyber threats. Attacks can cause unavailability or malfunction of systems which in turn prevent or mislead ongoing business processes in organizations. Today, organizations not only require cybersecurity programs to protect themselves against cyber threats but also need a resilience strategy to guarantee business continuity even during cyber incidents. This paper includes the results of ongoing research for securing maritime port ecosystems and making them cyber resilient. We propose a framework based on ontologies and logical inference to meet requirements of resilient IT systems regarding response to and recovery from cyber incidents.

1 INTRODUCTION

Hosseini et al. (2016) associate the term resilience to those of robustness, fault-tolerance, flexibility, survivability and agility. In this paper, we consider a system to be resilient, if it is able to operate during and to continue normally after an adverse event. Edwards (2009) depicts a resilience cycle in Figure 1 with five phases: Prepare, Prevent, Protect, Respond, and Recover. Cyber resilience defines the ability of an enterprise to effectively prepare against, prevent, detect, respond to, and recover from cyber incidents. If an enterprise is at least partially able to continue its business operations during a cyber incident, it will be called a cyber resilient enterprise.

The Prepare phase deals with the preparation against cyber attacks by designing and setting up vulnerability warning systems. The Prevent phase aims to implement measures for identified vulnerabilities to prevent cyber attacks as much as possible. As it is too good to be true to prevent all potential cyber attacks, some techniques are required to detect successful attacks. Threat detection is a very interesting subject of research in the area of cybersecurity and every year many new commercial products and academic papers are presented in the field. However, there is no product or framework that can guarantee the detection of all cyber incidents including zero-day attacks. The Respond phase deals with fast and effective emergency measures after an incident has been detected. These measures should compensate for damages done and take precautions to avoid further damages. The Respond phase aims to ensure the survival of the system. The Recover phase is the ultimate resilience goal of resetting the system into its original state or even into a better state in that lessons learned from past incidents are documented and contribute to future resilience.

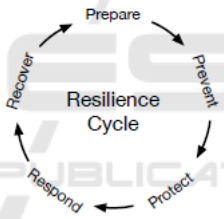

Figure 1: Resilience Cycle.

1.1 The Port Ecosystem

The maritime industry plays a major role in supplying consumers with overseas goods including raw materials such as energy, iron, and products in local retailers and supermarkets. In general, the shipping industry is responsible for 90% of trades in the world (ICS, 2020). The maritime supply chain is very dependent on port ecosystems with different stakeholders such as port operators, terminals, port authori-

Hutschenreuter, H., Çakmakçı, S., Maeder, C. and Kemmerich, T.
Ontology-based Cybersecurity and Resilience Framework.
DOI: 10.5220/0010233604580466
In *Proceedings of the 7th International Conference on Information Systems Security and Privacy (ICISSP 2021)*, pages 458-466
ISBN: 978-989-758-491-6

ties, navigation companies, shipping lines, and carrier organizations (cf. Figure 6). Stakeholders of the port ecosystem continue to adapt more Information Technology (IT) and Operational Technology (OT) systems to gain competitive advantages in the global maritime industry.

As the number of digital components like IT, OT, Internet of Things (IoT), and Bring Your Own Device (BYOD) is increasing, the port ecosystem becomes more vulnerable to cyber incidents and less cyber resilient. Unfortunately, the port ecosystem is still lagging behind in cybersecurity. Most of the port systems are not supported by state-of-the-art security solutions. Lack of effective maritime security regulation (Hopcraft and Martin, 2018), the conservative nature of the maritime industry, and lack of awareness can be seen as major challenges for establishing resilience. Surprisingly, there is a very limited number of academic studies in the field of cyber resilience in port systems.

1.2 Ontologies

In Computer Science an ontology is a formal description of concepts and relationships for an application domain of the real word (Staab and Studer, 2009). Computational ontologies are a means to formally model the structure of a system which are useful to our purposes. A central aspect is sharing of information and knowledge using a common vocabulary as supported by the *Resource Description Framework* (RDF). A family of ontology languages are *Description Logics* (DL) based on a well-understood subset of first-order predicate logics. The quest for expressive knowledge bases led to the development of the *Web Ontology Language* (OWL) as the current representation language of choice. The basic building blocks of OWL are concepts, roles, and individuals. Individuals present instances of concepts and roles relate them to each other. For classification purposes concepts are often hierarchically related to each other to form a taxonomy. So there may be super- and subconcepts, i.e. generalization and specialization of concepts. There are two flavors of relationships between instances of concepts that can be expressed via OWL: abstract and concrete roles. Abstract roles connect individuals and are also known as object properties. Concrete roles or data properties connect individuals with values of primitive data types (Pinkston et al., 2003). In DL and their semantics, the terms ABox and TBox are used to describe two different types of statements. TBoxes contain so-called *terminological schema knowledge* and ABoxes *assertional instance knowledge* (Hitzler et al., 2008). OWL does not provide a clear separation between TBoxes and ABoxes (Brockmans et al., 2004). Nevertheless, when describing OWL ontologies, it is helpful to introduce a separation between TBoxes and ABoxes which is based on intuitive criteria rather than being strictly formal. In OWL a TBox mainly describes static knowledge whereas an ABox is used for dynamic and frequently changing knowledge. Therefore, a TBox often infers, tracks or verifies class memberships while an ABox checks consistencies, facts, instances or operations in a rule-based manner.

1.3 Our Contribution

In this study, we present a novel resilience framework by combining a Security Information and Event Management Systems (SIEM) for detecting cyber attacks with ontologies and an inference system for supporting the selection of measures during cyber incidents and IT infrastructure failures in an automatic way. According to our literature review, there is no research in the academia and industry that addresses a cyber resilient port ecosystem framework. Of course, our proposed framework can be applied to provide cyber resilience for any IT-dependent organization.

The remainder of this paper is organized as follows. Section 2 presents a short literature review. Section 3 presents the details of the proposed framework. Section 4 describes a potential application of the proposed framework at a port terminal. Section 5 concludes with suggestions for future work.

2 LITERATURE REVIEW

Pinkston et al. (2003) and He et al. (2004) developed an ontology-based Intrusion Detection Systems (IDS) to spot different types of cyber attacks. Ekelhart et al. (2006) proposed a three parts ontology for supporting the cybersecurity of an organization including modeling knowledge about the security domain, existing IT infrastructure, and the people and their roles in the organization. Fenz et al. (2007) reported on a security ontology for the preparation of IT security audits that describe an ontological mapping of the structure of the IEC/ISO 27001 standard. Birkholz et al. (2011) described an ontology as a medium for the cross-company sharing of IT security knowledge. Elfers (2014) used an ontology to normalize events before event-correlation in the SIEM instead of storing the background knowledge in databases. Syed et al. (2016) defined a *Unified Cybersecurity Ontology* (UCO) to support the understanding of the underlying domain. UCO unifies the most commonly

used standards and enables the integration of public ontologies. It also enables sharing reasonable cybersecurity intelligence data. Petrenko and Makoveichuk (2017) developed an ontology of cybersecurity of self-recovering smart grids. Their devolved smart grid is resistant to negative impacts of the information confrontation and can quickly recover functions after accidents. Narayanan et al. (2018) developed a collaborative cognitive assistant based on ontology to detect cybersecurity events and attacks. Their proposed system collects incomplete textual information from different sources (e.g., CTI platforms) storing it in a structural format and reassign over it. Choi and Choi (2019) developed an ontology-based security context reasoning for attack detection in the power IoT-cloud environment. They proposed an attack context inference methodology based on the attack patterns in the power IoT-cloud environment and the vulnerabilities of each system in the major domains. Välja et al. (2020) developed an ontology framework to automatically model threats in the IT architecture of enterprises. Data is collected from enterprise resources (e.g., configuration data, captured traffic and scan data) that is parsed, standardized and merged. The reasoning patterns (i.e., queries on the ontology) are used for automated modeling of threats based on incoming data.

2.1 Maritime Cyber Attacks

In the past, cyber attacks have impacted individual stakeholders of the maritime industry like shipping companies and individual ports. Examples are the Danish Port Authority (2014), the Maersk shipping line (2017), ports of Antwerp (2013), San Diego (2018), Barcelona (2018), and Shahid Rajee (2020) in Iran. The complex and distributed characteristics of the port ecosystem makes it also vulnerable for a *distributed* cyber attack (Senarak, 2020).

2.2 Resilience and the Port Ecosystem

Recently, there is an increasing interest in cybersecurity and cyber resilience projects in Asian Pacific ports (Daffron et al., 2019), U.S. ports (LA, 2019), and EU ports (Rotterdam, 2016). The CyberShip project (Sepúlveda Estay, 2020) made a very detailed systematic literature review in the field of Cyber Resilience Frameworks (CRF). They selected five frameworks namely the Wave Analogy Model, the AWaRE framework, the Byzantine Fault tolerant framework, the Human Behavior Resilience framework, and the NIST framework. All frameworks were applied to ship operations and results were compared.

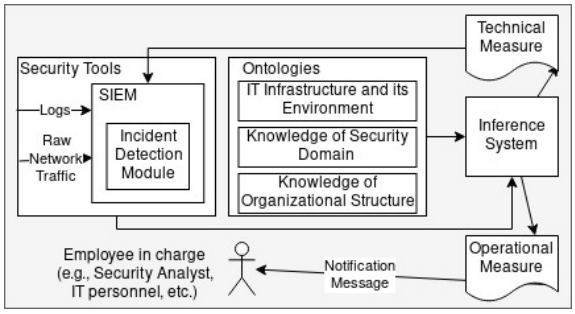

Figure 2: Resilience Framework.

3 PROPOSED FRAMEWORK

Our proposed framework is composed of three subcomponents as shown in Figure 2. The framework only uses ontologies and inference to response to and recover from cyber incidents in an automatic way. The inference system is not responsible for incident detection. Cyber incidents are detected by state-of-the-art and recent security tools.

The *security tools* of the resilience component should continuously monitor the IT systems using different kinds of sensors. Sensors collect specific and detailed information about what happens in the network of an organization and communicates it to the SIEM. Moreover, security tools like firewalls, antivirus, IDS, or Intrusion Prevention Systems (IPS) send their logs via Syslog messages over TCP or TLS to the SIEM, directly. The whole network traffic can also be captured by tools like Netflow (B. Claise, 2004) or packet sniffing (Fuentes and Kar, 2005) and fed to the SIEM for further analysis. Almost all SIEMs have capabilities to reduce noise, fine-tune alerts, and to identify inside or outside threats. Some recent SIEMs can also detect anomalies, i.e. they report unexpected behavior of any IT component including user input, hardware, and software. These anomalies need to be investigated by security analysts to be marked either as a dangerous activity or as a change of the normal behavior. For example, an increase in the number of requests for registering containers in a system might be a Distributed Denial of Service (DDoS) attack or just an increase in the trade rate. Berkovich and Solomon (2007) called an observed increase of trade before the new year a *Flash Event* (FE). The DDoS attack and the FE should be treated differently. This is where the inference system plays a significant role in our proposed framework by automatically applying required measures. Rule-based, correlation, statistical analysis, machine learning, deep learning are examples of currently used techniques to detect cyber incidents by SIEM.

3.1 Ontology of IT Infrastructures

The ontology of IT infrastructures serves to represent explicit knowledge about real IT systems and to enable reasoning based on this knowledge. We divide the infrastructure into two groups: Hard-IT and Soft-IT infrastructures. The latter is used by people to access stored data or use services available in Hard-IT infrastructures. Then, security measures will be applied to these IT infrastructures.

Figure 3 depicts the TBox (the static part of the ontology) for the Hard-IT. Solid arrows represent the relations between the classes as roles and inverse roles. For example, the class Network allows the modeling of networks using the subclasses Internet or Intranet. The role partOfNetwork is supposed to have a *functional characteristic* so that each network connection—an instance of the class Link—can only be assigned to exactly one network. If a physical network connection has to be assigned to several networks, e.g., by using Virtual LANs (VLAN), it has to be modeled by several parallel network connections. Other classes and roles of the Hard-IT infrastructure are to be interpreted in a similar manner.

For the practical usage of the ontology, an ABox, which is the dynamic part of the ontology, needs to be modeled in addition to the TBox. An ABox contains concrete instances of the classes and roles. For example, there may be two instances of the class PC (PC1 and PC2) and one instance of the class WiredLink (WLink1). Subsequently, the connectedToLink role should be used to describe the connection between PC1/PC2 and WLink1. From the ontology, it can now be inferred that PC1 and PC2 are connected to each other. In contrast to the TBox, the ABox of an ontology depends on the actual IT infrastructure of an organization in which our proposed framework will be applied. Summarizing, the TBox and an ABox together represent the following knowledge about an infrastructure:

1. The location of the Hard-IT components within the buildings and rooms.
2. The physical connection between different components of the Hard-IT (The IT devices are wired in a way that an appropriate configuration of them enables an actual network connection.)
3. The places where security measures can be applied on the Hard-IT level

Figure 4 displays the modeling of the TBox for the Soft-IT which is divided into Object and Subject classes. These classes are inspired by an Access Control (AC) model where subjects (entities, e.g., users, programs, processes) have certain permissions (e.g., rwx) to access objects (e.g., data, programs, devices) based on AC policies.

Elements of the Soft-IT infrastructure cannot exist outside the Hard-IT infrastructure. Therefore, all instances of the classes Subject and Object must be assigned to actual devices of the Hard-IT. The role validForDevice expresses that a subject is allowed to use a set of devices. Subjects such as users or groups can exist both locally for a device and globally within an administered computer network (e.g., with a directory service). Objects are assigned to devices via the role hostedOnDevice and their inverse role hasObject. Here again, hostedOnDevice is supposed to have a *functional characteristic* so that each object can only be assigned to exactly one device. If objects should be available on several devices, they must be modeled as copies.

In short, the TBox and an ABox represent the following knowledge about the Soft-IT infrastructure:

1. The possible accesses of subjects to objects
2. The places where security measures can be applied on the Soft-IT level

3.2 Ontology of Security

Figure 5 shows a TBox for cybersecurity. The bottom left depicts the TBox for the organizational structure via the class BusinessProcess to represent business processes, the class Role to represent organizational positions, and the class Employee to model employees who fill certain positions of a company.

Infrastructures are required for the execution of business processes. This is expressed by the role needsInfrastructure and its inverse role runsBusinessProcess. Infrastructures are indirectly accessed via roles that are assigned to responsible employees of an organization. The connection of responsible employees to infrastructures is given by the roles hasRole and accessesInfrastructure with corresponding reverse roles filledByEmployee and isControlledByRole. The following knowledge is represented by this part of the TBox and an ABox about a single organizational structure:

1. The responsible people who should be informed in the case of a cyber threat.
2. The responsible people who should carry out technical or organizational measures

The top parts of Figure 5 are the classes Threat, SecurityObjective and SecurityMeasure. Security measures are further subdivided into subclasses DetectMeasure, PreventMeasure, and ResonseRecoveryMeasure. The latter class for response and recovery measures further differentiates between *technical* and *organizational* measures.

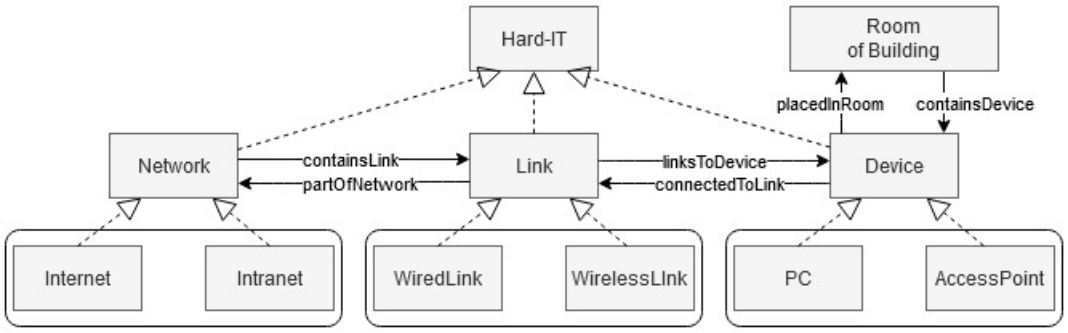

Figure 3: TBox for Hard-IT Infrastructures.

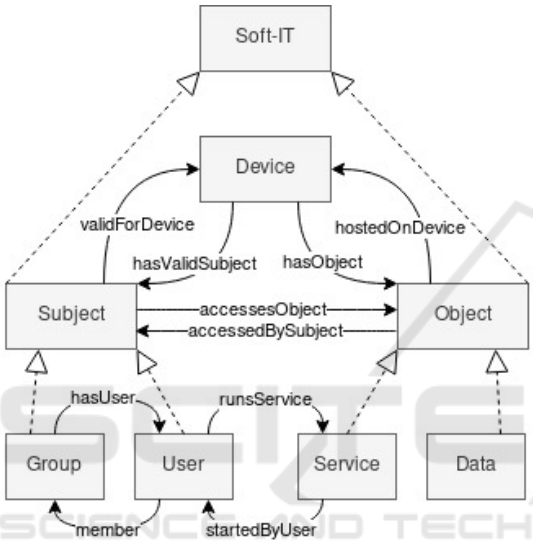

Figure 4: TBox for Soft-IT Infrastructures.

Security objectives are compromised by threats that in turn are reduced by security measures that support the security objectives. The relationships between security objectives, security measures, and threats exist independently of a specific application scenario. Therefore, individuals of these classes and their relationships can be modeled in an ABox independently from a specific infrastructure.

The security objectives should be determined by the protection requirements of an infrastructure. In our ontology, we have chosen to implicitly model the need for protection by relating the infrastructure to the security objectives. Thus, we simplified the security ontology without modeling the need for protection. Security measures are used to achieve the desired security objectives of an infrastructure. It is possible that several security measures support a common security objective, but some measures may only be suitable to protect a single infrastructure component. This is modeled by the role protectsInfrastructure and its inverse role isProtectedBySecurityMeasure. Actual technical measures are connected to infrastructure components by the role integratesTechnicalMeasure and its inverse role isIntegratedByInfrastructure. For example, a technical measure could be isolating the infected node in the network by blocking ingoing and outgoing traffic using a firewall. An additional organizational measures could ensure business continuity during an incident. Organizational measures are alternative solutions that affect the available business processes to ensure resilience. For example, when a website of an online shopping site is down, the company continues to receive orders via email or telephone.

3.3 Inference System

All developed ontologies should be stored to be accessible by interested agents or other system components. Subsequently, a reasoner is required to infer logical consequences (desired knowledge) from descriptive logics. Such a reasoner derives new statements from given statements and then extends the ontology with these new statements. For the proposed framework we need to use an OWL reasoner as an inference engine; e.g., HermiT (2013) that comes with Protégé (Musen, 2015). Designing an ontology in a manner that causes the inference engine to produce intended logical statements for its main purpose is a challenging task. It could be satisfied by a more complex ontology design or by simplified required conclusions. In the proposed framework, we aim to simplify the conclusions. For example, we do not want to determine emergency measures with a single run of the reasoner, instead we suggest using several runs and adjusting the ontologies between each run. As a result, we can reach a more complex conclusion via several simple conclusions.

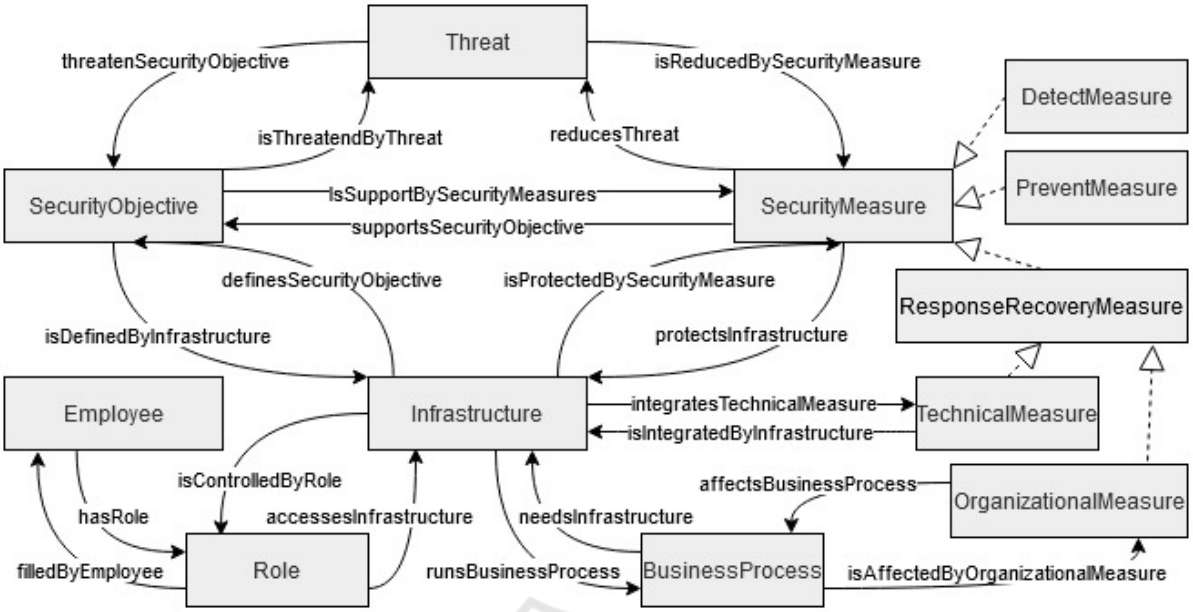

Figure 5: TBox for Cybersecurity.

4 FRAMEWORK APPLICATION

Each organization has different business objectives, business processes, and infrastructures. Therefore, each organization requires unique modeling of its infrastructure and security to protect the continuity of its main business, i.e. to keep executing its business processes while facing security threats.

4.1 Case Study: Stowage Planning

This paper demonstrates a potential implementation of our proposed solution using an example of a Port Community System (PCS). A PCS is a central part of a complex and distributed system in which different stakeholders such as terminal operators, ship owners, forwarders, port IT operators, railways, port authorities and customs network with each other. Figure 6 shows a simplified representation of a network around a PCS.

IT plays a central role for communication between the different stakeholders. To keep port communication systems resilient, all stakeholders should take care of their internal cyber resilience. There is a large number of processes within a port community. For demonstration purposes, we only select a single subprocess which is mainly handled by the terminal. Figure 7 shows the stowage planning subprocess at the port terminal as a sequence diagram. The PCS forwards the customs clearance message to the terminal.

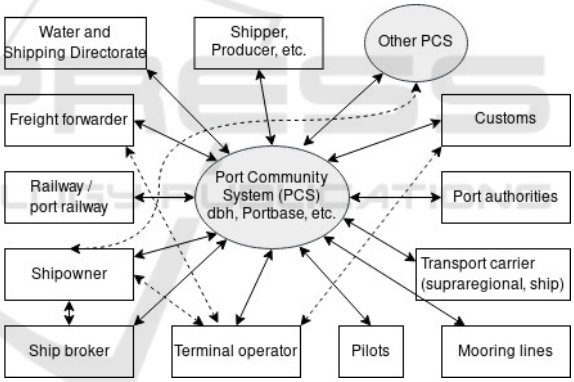

Figure 6: Port Community System.

Goods must not be loaded without customs clearance. The terminal and the vessel operator negotiate the unloading and loading of goods from and to the vessel via a stowage- or bayplan, also known and abbreviated as BAPLIE. This document contains information such as the position of containers on board and the weight of the containers. Since 2016, containers in the EU require a *Verified Gross Mass* (VGM) to increase safety on sea. Moreover, this document includes special information about dangerous goods or cooling requirements for temperature-sensitive cargo.

4.2 Resilience Framework in Use

Figure 8 depicts potential IT equipment and our proposed resilience framework of an imaginary port ter-

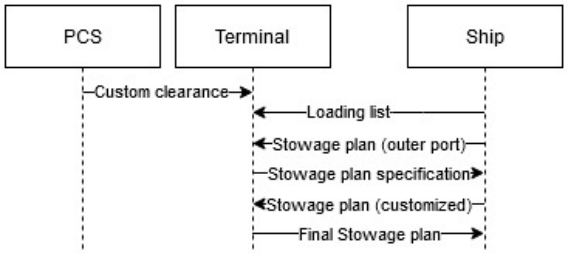

Figure 7: Stowage Planning Process.

minal in a simplified way. All IT infrastructures are supposed to send their logs for storage and analysis to the SIEM. When the SIEM detects an abnormal event or a specific attack, it sends an alarm to the inference system. The communication between SIEM and the inference system is based on details of the ontologies. In this study, we use low-level models which imply low-level communications between the SIEM and the inference system. For our example, the SIEM only needs to provide two pieces of information to the inference system. Whenever an incident is detected we need to know the compromised/damaged IT *infrastructure* (e.g., which PC, router, etc.) and the attack/anomaly *type* (e.g., failure, malware, DDoS). Subsequently, the inference system should take the following steps:

1. The attack is transferred into the ontology by adding facts about compromised or damaged infrastructure components (Hard-IT, Soft-IT).
2. The inference system starts a reasoning process and infers the full impact of a failure including further potential threats and affected business processes.
3. The inference system provides some suggestions of possible measures to overcome the failure, as threats and measures are related to each other by the security ontology.
4. To evaluate the effectiveness of a measure, the measure is temporarily added as fact to the ontology.
5. The inference system performs another reasoning step to check if the situation has improved in which case the measure is kept as fact and conveyed to the SIEM.
6. Steps 4 to 5 are repeated as long as the situation improves and no further risks threaten business processes.

4.3 Attack Scenarios with Automatic Response and Recovery

A SIEM detects a known cyber incident or abnormal behavior of a Terminal Operating Server (TOS) or communicating workstations and reports incidents to the inference system. The inference system checks its knowledge-base for automatic responses. The following list shows some of the potential SIEM detection messages and corresponding response and recovery proposals of the inference system:

Malware Is Detected on TOS. Install an updated antimalware tool, rescan the TOS server
Malware on TOS Is Not Removed by Antimalware. Disconnect the TOS server, create a backup of files, deploy a redundant TOS server, inform responsible IT personnel to eliminate the malware manually
Ransomware Is Detected on TOS. Disconnect the TOS server, wipe the system, install from scratch, restore backup files, inform responsible people, inform the authorities, attempt to decrypt files by trusted tools, deploy a redundant TOS server
DDoS Is Detected on TOS. Increase bandwidth on the TOS server, locate proxy servers in front of the TOS server with a load balancer, distribute arriving requests between proxy servers via the load balancer, enlarge the backlog queue for requests to the TOS, increase the timeout of connections in backlog, block suspicious Geo-IP addresses, inform responsible people, inform authorities
Abnormal Behavior on TOS. Make backup, analyze logs, inform responsible people
Abnormal Activities on a Workstation Communicating with the TOS. Disconnect the workstation and its LAN, collect and analyze data from the workstation (Application logs and system parameters), assign a new workstation for communicating with the TOS, inform responsible people
TOS Is Offline for More than One Hour. Inform responsible people, inform partners who will be affected by the unavailability of TOS, make announcement on the webpage of terminal
TOS Is Recovered by IT-staff. Unplug a redundant TOS server, deploy the main TOS server, restore backup files, re-inform partners, update announcement from the webpage of terminal

The inference system provides responses based on its knowledge about available measures, IT infrastructures, and security objectives. For example, if there is no redundant TOS server in the terminal, the inference system will send an Inform-Partners message rather than proposing to deploy the redundant TOS server. If the terminal does not have any load balancer between its IT infrastructures, the inference system does not select a measure that requires the installation of load balancer.

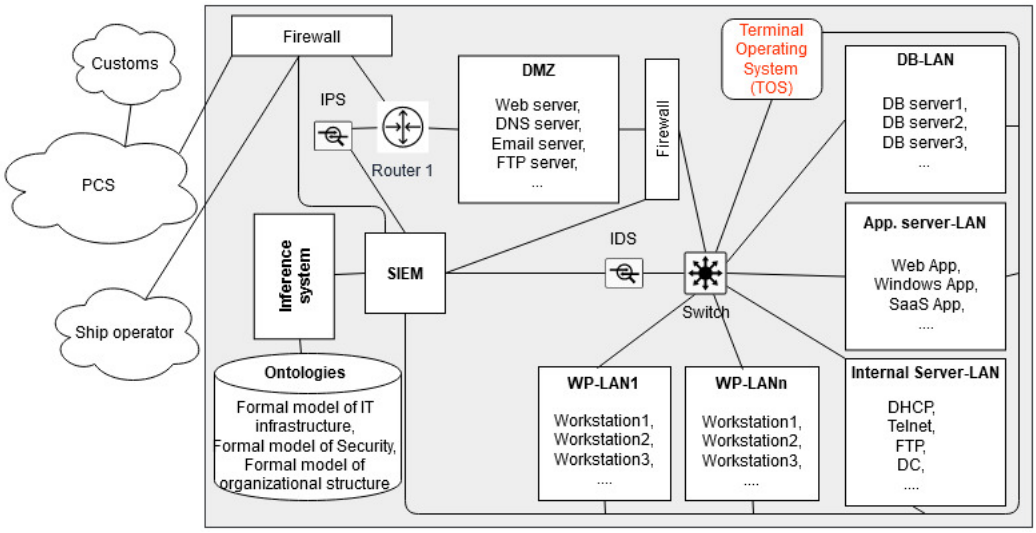

Figure 8: A high-level overview of an imaginary terminal's network.

Nowadays, most modern SIEMs can act automatically including isolating, disabling, removing, or changing Hard-IT infrastructures and their configurations (Visser, 2020). Therefore, we assume that some scripts should be newly written for terminal SIEMs to execute special commands when receiving messages from the inference system. For example, if a SIEM receives a Disconnect-TOS message, it will reconfigure the router setting to isolate the TOS server and prevent network traffic from and to the TOS server.

5 CONCLUSIONS

Today, cyber resilience should work in alliance with traditional cybersecurity to protect organizations against cyber incidents. Cyber resilience is more concerned with respond to and recovery from cyber incidents. There are many traditional and novel tools like antivirus, firewall and second-generation SIEMs to prevent and detect incidents. Nevertheless, incident response and recovery are often manual processes done by security experts. In this paper, we proposed a novel approach for the application of semantic technologies. Our ontologies include the most basic definitions of IT infrastructures and security. Each of these models can be expanded and fine-tuned based on the requirements and resources of an organization. This paper is based on the result of an ongoing project in the maritime transport industry. Port ecosystems are far behind other industries regarding cybersecurity as the number of available experts seems to be rather limited. We try to develop a new framework which automates not only the detection of cyber incidents but also the response and recovery phases. Subsequently, lessons learned from cyber incidents and taken measures should also improve reoccuring preparation, prevention, and protection phases of the resilience cycle. The current paper lacks empirical results that should be addressed in future work. The detailed modeling of ABoxes will be based on the organizational and technological structure of a concrete terminal. Calculations of IT infrastructure *capacities* during incidents and exchanging these with business partners is another interesting area for future work in the field of cyber resilience.

ACKNOWLEDGMENTS

This work was supported by the German Federal Ministry of Transport and Digital Infrastructure (BMVI) under the grant 19H18012E (SecProPort project).

REFERENCES

B. Claise, E. (2004). Cisco systems NetFlow services export version 9. https://www.hjp.at/doc/rfc/rfc3954.html.

Berkovich, E. and Solomon, O. (2007). Flash event detection with acoustic verification. US Patent 7,233,546 https://patentimages.storage.googleapis.com/00/e8/1b/100810723fac4e/US7233546.pdf.

Birkholz, H., Elfers, C., Samjeske, B., and Sohr, K. (2011). Unternehmensübergreifender Austausch von sicherheitsrelevantem Wissen. *Datenschutz und Datensicherheit-DuD*, 35(4):258–261. doi:10.1007/s11623-011-0063-5.

Brockmans, S., Volz, R., Eberhart, A., and Löffler, P. (2004). Visual modeling of OWL DL ontologies using UML. In *International Semantic Web Conference*, pages 198–213. Springer. doi:10.1007/978-3-540-30475-3_15.

Choi, C. and Choi, J. (2019). Ontology-based security context reasoning for power IoT-cloud security service. *IEEE Access*, 7:110510–110517. doi:10.1109/ACCESS.2019.2933859.

Daffron, J., Ruffle, S., Coburn, A., Copic, J., Quantrill, K., Strong, K., and Leverett, E. (2019). Shen attack: Cyber risk in asia pacific ports. https://www.lloyds.com/~/media/files/news-and-insight/risk-insight/2019/shen-attack/cyrim_shenattack_finalreport.pdf.

Edwards, C. (2009). *Resilient nation*. Demos. https://www.continuitycentral.com/ResilientNation.pdf.

Ekelhart, A., Fenz, S., Klemen, M. D., and Weippl, E. R. (2006). Security ontology: Simulating threats to corporate assets. In *International Conference on Information Systems Security*, pages 249–259. Springer. doi:10.1007/11961635_17.

Elfers, C. (2014). *Event Correlation Using Conditional Exponential Models with Tolerant Pattern Matching Applied to Incident Detection*. Shaker. https://www.shaker.de/shop/978-3-8440-3168-3.

Fenz, S., Goluch, G., Ekelhart, A., Riedl, B., and Weippl, E. (2007). Information security fortification by ontological mapping of the iso/iec 27001 standard. In *13th Pacific Rim International Symposium on Dependable Computing (PRDC 2007)*, pages 381–388. IEEE. doi:10.1109/PRDC.2007.29.

Fuentes, F. and Kar, D. C. (2005). Ethereal vs. tcpdump: a comparative study on packet sniffing tools for educational purpose. *Journal of Computing Sciences in Colleges*, 20(4):169–176. https://dl.acm.org/doi/10.5555/1047846.1047873.

He, Y., Chen, W., Yang, M., and Peng, W. (2004). Ontology based cooperative intrusion detection system. In *IFIP International Conference on Network and Parallel Computing*, pages 419–426. Springer. doi:10.1007/978-3-540-30141-7_59.

HermiT (2013). HermiT OWL Reasoner. http://www.hermit-reasoner.com.

Hitzler, P., Krötzsch, M., Rudolph, S., and Sure, Y. (2008). *Semantic Web: Grundlagen*. Springer. doi:10.1007/978-3-540-33994-6.

Hopcraft, R. and Martin, K. M. (2018). Effective maritime cybersecurity regulation–the case for a cyber code. *Journal of the Indian Ocean Region*, 14(3):354–366. doi:10.1080/19480881.2018.1519056.

Hosseini, S., Barker, K., and Ramirez-Marquez, J. E. (2016). A review of definitions and measures of system resilience. *Reliability Engineering & System Safety*, 145:47–61. doi:10.1016/j.ress.2015.08.006.

ICS (2020). Shipping and world trade. https://www.ics-shipping.org/shipping-facts/shipping-and-world-trade.

LA (2019). Port of Los Angeles issues request for proposal for new cyber resilience center. https://www.portoflosangeles.org/references/news_072419_rfp_cyber_resilience_center.

Musen, M. A. (2015). The Protégé project: A look back and a look forward. *AI Matters*, 1(4):4–12. doi:10.1145/2757001.2757003.

Narayanan, S. N., Ganesan, A., Joshi, K., Oates, T., Joshi, A., and Finin, T. (2018). Early detection of cybersecurity threats using collaborative cognition. In *2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC)*, pages 354–363. doi:10.1109/CIC.2018.00054.

Petrenko, S. A. and Makoveichuk, K. A. (2017). Ontology of cyber security of self-recovering smart grid. In *Secure Information Technologies (BIT 2017)*, pages 98–106. http://ceur-ws.org/Vol-2081/paper21.pdf.

Pinkston, J., Undercoffer, J., Joshi, A., and Finin, T. (2003). A target-centric ontology for intrusion detection. In *Workshop on Ontologies in Distributed Systems, held at The 18th International Joint Conference on Artificial Intelligence*. https://ebiquity.umbc.edu/get/a/publication/626.pdf.

Rotterdam (2016). Port of Rotterdam appoints port cyber resilience officer. https://www.portofrotterdam.com/en/news-and-press-releases/port-of-rotterdam-appoints-port-cyber-resilience-officer.

Senarak, C. (2020). Port cybersecurity and threat: A structural model for prevention and policy development. *The Asian Journal of Shipping and Logistics*. doi:10.1016/j.ajsl.2020.05.001.

Sepúlveda Estay, D. A. (2020). *CyberShip Project: Cyber resilience for the shipping industry - Final Project Report*. DTU Orbit. https://orbit.dtu.dk/files/216595381/200630_Report_WP_5.pdf.

Staab, S. and Studer, R. (2009). *Handbook on Ontologies*. Springer. doi:10.1007/978-3-540-92673-3.

Syed, Z., Padia, A., Finin, T., Mathews, L., and Joshi, A. (2016). UCO: A unified cybersecurity ontology. In *Proceedings of the AAAI Workshop on Artificial Intelligence for Cyber Security*. AAAI Press. http://ebiquity.umbc.edu/get/a/publication/781.pdf.

Välja, M., Heiding, F., Franke, U., and Lagerström, R. (2020). Automating threat modeling using an ontology framework: Validated with data from critical infrastructures. *Cybersecurity*, 3(19). doi:10.1186/s42400-020-00060-8.

Visser, J. (2020). An OODA-driven SOC Strategy using: SIEM, SOAR and EDR. http://correlatedsecurity.com/an-ooda-driven-soc-strategy-using-siem-soar-edr.