EUROPEAN UNION AGENCY
FOR CYBERSECURITY

ARTIFICIAL INTELLIGENCE

Conference Report

SEPTEMBER 2019

ABOUT ENISA

The mission of the European Union Agency for Cybersecurity (ENISA) is to achieve a high common level of cybersecurity across the Union, by actively supporting Member States, Union institutions, bodies, offices and agencies in improving cybersecurity. We contribute to policy development and implementation, support capacity building and preparedness, facilitate operational cooperation at Union level, enhance the trustworthiness of ICT products, services and processes by rolling out cybersecurity certification schemes, enable knowledge sharing, research, innovation and awareness building, whilst developing cross-border communities. Our goal is to strengthen trust in the connected economy, boost resilience of the Union's infrastructure and services and keep our society cyber secure. More information about ENISA and its work can be found at www.enisa.europa.eu.

CONTACT

For contacting the authors please use COD3@enisa.europa.eu.
For media enquiries about this paper, please use press@enisa.europa.eu.

AUTHORS

Georgios Chatzichristos, Fabio di Franco, Ioannis Agrafiotis, Cosmin Ciobanu

ACKNOWLEDGEMENTS

ENISA would like to thank all speakers, panellists and contributors to this conference for their valuable inputs that this report summarizes.

LEGAL NOTICE

Notice must be taken that this publication represents the views and interpretations of ENISA, unless stated otherwise. This publication should not be construed to be a legal action of ENISA or the ENISA bodies unless adopted pursuant to the Regulation (EU) No 2019/881.
This publication does not necessarily represent state-of the-art and ENISA may update it from time to time.

Third-party sources are quoted as appropriate. ENISA is not responsible for the content of the external sources including external websites referenced in this publication.

This publication is intended for information purposes only. It must be accessible free of charge. Neither ENISA nor any person acting on its behalf is responsible for the use that might be made of the information contained in this publication.

COPYRIGHT NOTICE

TABLE OF CONTENTS

EXECUTIVE SUMMARY

On the 3rd and 4th of June 2019, ENISA organized in Athens, Greece, a conference on Artificial Intelligence in the EU Cyber crisis blueprint context under the title ‘ Artificial Intelligence-An opportunity for the EU cyber-crisis management’.

Artificial Intelligence (AI) is not something new. What is new is the pace that information is disseminated in today’s society and the amount of data produced that renders the use of advanced technologies like AI essential for organisations. AI technology, if properly applied, has the potential to create a competitive advantage over traditional methods of data classification. An increasing number of public and private organisations are using now AI, both at operational and technical levels. AI systems are utilised to provide intelligence on the services which organisations offer, to distinguish behavioural aspects of systems and networks, and to help humans understand complex relationships between different entities of their working environment. In cybersecurity, AI systems have the ability to highlight anomalies on network traffic identifying invisible “unknown unknown things” in the systems, but also as an efficient classifier of vast amounts of data like in the case of threat intelligence.

In the EU cyber crisis cooperation context, most commonly referred as ‘the Blueprint’, AI uses are just beginning to emerge, mainly at the Open Source intelligence domain. The need for situational awareness, one of the Blueprint’s main pillars, has driven ENISA to initiate the development of a project under the name ‘Open Cyber Situational Awareness Machine – OpenCSAM) that attempts to address the need for accurate aggregation of relevant information and reporting. The project is using supervised learning and natural language processing to facilitate incident responders at all levels of administration in the drafting of situational awareness reports for the Blueprint.

This report summarizes the main takeaways from the two days of the conference and the discussions of the four (4) thematic panels from the Political to the Technical level. The presentations that were delivered in the context of the four (4) thematic areas can be found here: https://www.enisa.europa.eu/events/artificial-intelligence-an-opportunity-for-the-eu-cyber-crisis-management/workshop-presentations/event-presentations

1. CONFERENCE ORGANIZATION

1.1 SCOPE OF THE CONFERENCE

The conference aimed at supporting ENISA's activities for cyber-crisis management[1]. Another important driver has been the investigation of ways that Artificial Intelligence can be used in the Blueprint context to improve cyber security and to assist decision making.

The event was organized in two days and in four different sessions with corresponding discussion panels. In the first day, the centre of gravity was on the Political and Operational collaboration and the possible uses of Artificial Intelligence in these two levels of governance. In the following day the focus shifted to technical level discussions, AI solutions and applications already in use, challenges and future trends. There has been also a dedicated session to ENISA's Open Cyber Situational Awareness Machine (OpenCSAM), which is a research and development project that relies heavily on artificial intelligence and aims at supporting the reporting functions of the Blueprint.

1.2 PARTICIPATION

ENISA attempted with this event to bring together cyber and AI stakeholders across different sectors and domains. In the speakers' list, there were representatives from EU Institutions, the European Commission, The European External Action Service, EDA, Europol EC3 and Cert EU while among the audience there were representatives from the European Council, Eurocontrol and the ESDC. Academia had also a very strong presence with speakers from Academic Institutions well known for their research in cyber security and artificial intelligence. Finally yet importantly, public authorities and private companies delivered their inputs and presented solutions for the autonomous cyber response area.

Overall, over 120 participants from 12 Member States were present during the event. It is important to mention that almost half of the participants were either lawyers or related to legal and legislative activities in the public and in the private sector.

1.3 THEMATIC AREAS

The conference covered four thematic sections. Sections 1 and 2 during the first day focused on the Blueprint, Section 3 on the second day focused on ENISA's OpenCSAM project and the last section focused on autonomous cyber response at the technical level.

[1] ENISA Work Program 2019 Outputs 4.1.3 and 4.1.4

2. AREA 1 – THE FUTURE OF THE BLUEPRINT

The aim of the first panel was to verify the current status of cooperation between the key EU cyber stakeholders and discuss on the enhancement and possible expansion of this collaboration. In the panel there were representatives from the European Commission/DG CNECT (Dr. Ioannis Askoxylakis), the European External Action Service (Ms Agniezka Wierzbicka), The European Defence Agency (Mr Mario Beccia), Cert EU (Mr Georgios Psykakos), Europol EC3 (Ms Aglika Klayn) and the discussion was moderated by ENISA (Mr Georgios Chatzichristos).

2.1 GAPS AND TRENDS FOR THE BLUEPRINT FUTURE

The Commission mentioned that there is a big push for further Operationalization of the Blueprint. EU cyber exercises in the last years are revealing the need. From the Member States side, the NIS Cooperation Group established under the NIS Directive is active working with the Commission on the implementation of the Blueprint at the Operational level of the Member States' mechanisms. The four (4) MoU stakeholders (ENISA, EUROPOL EC3, EDA and CertEU) are the 'right' players as they represent all critical functions of cyber security in the EU ecosystem. However as the impact of cyber-attacks is not only on those players, all other EU Institutions and bodies should not only have a certain level of preparedness but also a role in the Blueprint as they could be called to exchange information, produce impact assessments, and respond to incidents on their domain of activities. In addition, due to the borderless nature of cyber space, the need for cooperation with non EU Countries that share the same core values is of key importance for the EU.

2.2 AN EFFECTIVE AND HOLISTIC APPROACH

Europe should avoid building new mechanisms, policies and capabilities without taking into account two important factors.

Utilize existing crisis management mechanisms, protocols, policies and experience: The Commission said that the Blueprint goal is to align current crisis management mechanisms and use good practices from other Sectors (eg. Banking Sector or Aviation Sector) regarding crisis management. EUROPOL EC3 added that we should borrow experience from existing work, such as the LE ERP protocol.

Ensure effectiveness: The measurement of effectiveness has been identified as of crucial importance for the Blueprint. EDA emphasized this fact, agreeing with the University of Oxford that this should be an area of research for the EU.

2.3 BUILDING TRUST

Cyber collaboration needs to address also sensitivities and reluctance in collaborating from some National Authorities. EDA gave some analogies in the capability development from its own experience in the Common Security and Defence Policy (CSDP) domain. Some Member States might be willing more to collaborate (mainly smaller Member States) while others are more reluctant. The EEAS added that in their experience for overcoming such problems, capacity building is important. This will lead to an increase of the maturity level and will lead to better collaboration. EUROPOL EC3 strongly agreed on this. This expands even beyond Member States, into third countries. Trust building is also of great importance. Trust is a continuous process that requires continuous efforts to build. We do not trust Organizations and mechanisms but we do trust the humans that make them work. EDA said that the Banking sector, for example, uses many circles or communities of people who come often together, get to know each other and build trust. Before having protocols, capabilities and information exchange networks, we should invest in the human factor first. CERT EU complemented that day-to-day Operations also lead to human interaction with different entities and this help to build trust. The ESDC presented another aspect of trust by connecting it to the attribution issue. Attribution, very often relies on classified Intelligence information that cannot be disclosed to all parties of the incident response mechanisms or to the public. In that way, incident response mechanisms are vulnerable to targeted propaganda attacks. EUROPOL EC3 emphasized the need of transparency in the processes. EEAS said that attribution is primarily a responsibility of the Member States. Attribution based on National Intelligence sources should be very carefully handled, especially regarding the communication of public messages, balancing between the protection of National Security interests and the need for transparency, a crucial factor for trust. ENISA said that for trust building in incident response we need to take into account three elements. '*Who am I, Who are you and what will you do with the data that I am giving you'*. A balance in the information exchange is also critical in the sense that '*I will give you something that you need and you will give me something that I need.*'

2.4 CIV-MIL COLLABORATION

Today's APTs and hybrid threats prove that that the Cyber domain is increasingly used for complicated cyber operations that target both civilian and military targets. In 2016, NATO declared the cyber domain as a domain of military operations[2]. This brought a revision in operational concepts and planning and introduced new legal concerns[3]. The cyber domain differs from the traditional Operational domains (Air, Sea, Land and Space) in the sense that is being depended and operated by the industry (i.e. telecom companies, IT manufacturers, IT service providers, media etc.). Because cyber domain's interdependencies are so complex and borders between military and civilian uses are blurry or non-existent, the effective civilian – military collaboration is so important. The Blueprint should act as a bridge that will bring closer the Civilian and the CSDP worlds for the better mitigation of common cyber threats. Both EEAS and EDA agreed the civil – military cooperation has always been a very complicated case. It is therefore next to impossible for the military not to take into account the civilian world when it comes to cyber security. Joint exercises is a good first step towards this collaboration.

2.5 CONCLUDING REMARKS

- Foresee roles and procedures for all EU Institutions in the Blueprint context.
- Need for continuous operations of certain blueprint functions, like Situational Awareness and Horizon scanning
- Build trust by developing information exchange networks, personnel exchange and frequent exercises between Blueprint stakeholders at the Operational level.
- Identify synergies between the Civil and Military domains in information exchange and capabilities.

[2] https://www.nato.int/cps/en/natohq/official_texts_133169.htm
[3] https://ccdcoe.org/research/tallinn-manual/

3. AREA 2 – HOW CAN AI HELP THE BLUEPRINT

The aim of the second panel was to discuss developments in the field of Artificial Intelligence and explore how advances in this area can be utilised in the implementation of the Blueprint. The panel was comprised of participants with diverse backgrounds, working in public and private institutions as well as in academia. More specifically, there were representatives from Ministry of Defense (MoD) of Bulgaria (Dr George Sharkov), the European Defense Agency (Mr Mario Beccia), Microsoft (Mr Maciej Surowiec), University of Oxford (Dr Jassim Happa) and Europol (Mr Dimitrios Zacharis). The panel was moderated by ENISA (Dr Ioannis Agrafiotis).

3.1 THE NEED FOR AI IN ORGANIZATIONS AND ITS PERILS

The pace with which information is disseminated in today's society and the amount of data produced renders the use of advanced technologies essential for organisations. AI systems have the potential to create competitive advantage and have been adopted by organisations both at operational and technical levels. Such systems are utilised to provide intelligence on the services, which organisations offer, to distinguish behavioural aspects of systems and networks, and to help humans understand complex relationships between different entities of their working environment. In cybersecurity, AI systems have the ability to highlight anomalies on network traffic, indicating possible cyber-attacks and facilitate analysis to identify invisible "unknown unknown things" in the systems.

The adoption of AI systems, however, has its perils. Participants concurred that recent advances in networking and storage infrastructure have enabled the implementation of AI algorithms that were invented in the 1970s. We have not yet realised the full potential of AI systems and the effects that these could have on our society. Concerns were raised on how AI systems are trained, since there is evidence to suggest that data poisoning of training datasets can manipulate the performance of AI systems. Furthermore, participants explained that widely applied algorithms, such as neural networks, act as black boxes, reducing transparency in how their outcome is produced. Such opaque behaviour can further introduce ethical dilemmas and biased decisions. Finally, participants explained that AI systems can be utilised by adversaries to design malware.

3.2 AI IN THE CONTEXT OF CYBERSECURITY AND THE BLUEPRINT

Participants suggested that organisations mainly use AI techniques in cybersecurity to detect anomalies in their network environments. A possible explanation given by the participants for such wide adoption of anomaly detection systems is that "it is more costly to deploy systems that protect against attacks rather than implement solutions which detect or mitigate successful attacks". There is a plethora of products offering behavioural analysis, however, such solutions operate in silos and it is becoming increasingly difficult to correlate data from such sources for better situational awareness. It was further mentioned, that solutions claim to utilise AI to detect 0-day attacks, however, participants questioned the validity of such tools especially when you need a 0-day attack deployed in a network to test the accuracy of such systems.

Power asymmetries between attackers and defenders was another topic of discussion where AI can potentially change the balance. Attackers have no ethical constraints and with risk appetite to try novel technologies. On the other hand, people who defend networks for organisations abide to specific regulations and new technologies have to be fully tested, justified and economically efficient before being deployed. Participants deemed that with AI systems, defenders can process vast amount of data, unveil correlations and make informed decisions faster, all of which can help to bridge the gap in power asymmetries and potentially reduce significantly the amount of time needed to detect Advanced Persistent Threats (APTs) (according to panellists, organisations in average require 150 days to identify APTs in their networks).

In order to deploy AI systems successfully, participants explained that organisations should expose these in a controlled environment first and gradually release them on the rest of the network. It is of paramount importance to have analysts training these systems to ensure that data from potentially compromised networks will be properly handled and will not compromise the output of AI systems. Participants further mentioned that the notion of redundancy is important when such systems are adopted and organisations should consider and combine inputs from different AI algorithms for better-informed decisions. Finally, benchmarking of such tools is another step that the security community should work towards, since most systems are treated as black boxes due to Intellectual Property (IP) issues (manufacturers of AI systems are not transparent in the algorithms as well as the features these systems use). Therefore, benchmarking can help organisations to identify which AI systems in which context are more accurate. Participants suggested that these are all valuable lessons which should be reflected in the AI system for Blueprint.

3.3 ACCOUNTABILITY OF AI AND ASSURANCE OF QUALITY

To address problematic situations from the use of AI, participants discussed the issue of accountability and quality assurance for AI systems. AI systems in essence utilise optimisation techniques to classify events and detect anomalies. Therefore, such systems cannot guarantee the quality of the outcome, especially those systems that depend on neural network techniques, which are not transparent in how decisions are made.

Participants suggested that we need to systematically test AI systems and learn how these perform in different contexts and how reliable their results can be. For black box systems, we need to obtain visibility in which neurons are ‘excited’ on given information and start building a picture on which features influence decisions. Furthermore, emphasis should be given on standardising the datasets based on which systems are trained. Ethical guidelines and potential biases should be taken into account when constructing such datasets. Finally, participants concurred that as with any other technology, organisations will adapt to the use of AI systems in cybersecurity and trust in such systems will be fostered through experience.

3.4 FAKE NEWS AND THE BLUEPRINT

The system that will support the implementation of the Blueprint will consume information from different sources with varying degrees of trustworthiness. Therefore, AI techniques could be utilised to help users fuse all this information appropriately. Particular emphasis was given on how misinformation and disinformation attacks, which can compromise the trustworthiness of the intelligence created by combining information from different sources, can be detected and mitigated.

Participants suggested that AI algorithms can cluster automatically trustworthy sources by identifying how information is disseminated, from which sources, which rumours are corroborated by which sources and what is the network topology of rumours. AI systems can then provide warnings of potentially untrusted sources to users before fusing information. Finally, participants suggested that in order to tackle the problem of fake news holistically, we need to compliment AI systems with a cultural change in users’ mindset. Therefore, they suggested that it is of paramount importance to educate users to think critically when reading and reproduce information.

3.5 CONCLUDING REMARKS

- Gradually expand the implementation of AI systems, from a controlled environment to the rest of the network.
- Carefully select the training dataset for the algorithms.
- Involve humans in the training of AI systems to ensure that quality of outputs is acceptable.
- Educate users on critical thinking and on evaluating the trustworthiness of sources.
- Collaborate with other stakeholders like the European Commission’s high level expert group on AI on the development of guidelines and assessment tools on trustworthy artificial intelligence

4. AREA 3 – IMPROVING OPENCSAM

The aim of the 3rd Panel was to discuss the technical challenges for the OpenCSAM project and identify future work for the next cycles of development. The panel was comprised of participants with diverse backgrounds, working in EU institutions, the private sector as well as in academia. More specifically, there were representatives from SDL (Mr George Bara), the University of Oxford (Dr Jassim Happa), ESDC (Dr. Gregor Schaffrath), ENISA's contractor for OpenCSAM - Eau de Web (Mr V.Posea & Mr T.Ichim), and the BCU School of computing and digital technology (Dr. Syed Naqvi). The panel was moderated by ENISA (Mr Cosmin Ciobanu). The panel discussion kicked off with the background and goals for developing the OpenCSAM tool. Being a tool developed for the blueprint, the decision makers require tailored situational awareness reports that provide context about ongoing developments during a cybersecurity incident.

4.1 EXTRACTION OF INFORMATION

One major technical difficulty in the development is the extraction of information (text) from the sources that OpenCSAM is using. What is the state-of-the-art on this matter, how it can be improved, to what extent you can automate the scraping and text extraction?

To this question the experts agreed that indeed this is a complex problem, incurring high costs for scaling, technically complicated and challenging due to heterogeneous web technologies. There are some platforms that can do the job with reasonable success. Most of the websites are dynamic nowadays and this might create difficulties. The best way to scrape text from websites is by using computer vision/image & character recognition techniques. We don't need to store everything (all the text from all the websites) we are interested in the cyber security related topics which could be simplify the problem. One essential aspect noted by the panellists is "from where should we scrape data?", from which sources, which are the "best/optimal" sources. Where are the pockets of information that we are really interested in and what metrics we can define for this purpose?. The knowledge graph is a curated ontology that helps in identify the right kind of information.

4.2 TEXT, TOPIC, LANGUAGE AND SENTIMENT CLASSIFICATION

The next topic of discussion was on text, topics, language and sentiment classification. Again here, the question posed was what is the state-of-the-art on this matter. It was mentioned that Google has released a generic model for text classification highly dependent on computing power. This can be fine-tuned for specific purposes, and it can be trained for specific use cases. The biggest challenge is having labelled data, essentially batches of text that is "known" to belong to a certain category. Another issue highlighted by the panellists relates to lack of labelled data for different languages other than English. Uncertainty and Trustworthiness were raised as a potential issue - to what degree the information collected reflects the reality. Coupled with this problem is also the issue of semantics and meaning of specific terms which could interpreted differently nowadays (different connotations) as compared to their initial meaning in the past with impact false positives rates. The meaning of a term should always be kept within the context of the article or resource.

4.3 COMPANY-ENTITY EXTRACTION

The next challenge that was discussed was the extent to which it is possible to extract the different entities from a text. Assuming that we have a cyber-attack article for example, can we identify the “parties” involved in the story? This is a classical task in Natural Language Processing (NLP) code-named ‘entity recognition’, and there are some quite advanced tools that can be used together with specific datasets with good results. Specific entities need to be defined, using open-source tool ‘Prat’. The analyst will have the task of training the algorithms and build upon a corpus of specific entities relevant to the specific domain. Besides algorithm training but also training the users in understanding the probabilistic nature of the systems that they are using is important. An interesting point raised was if we know what is the percentage of the data and how accurately a text is translated by for ex. Google Translate? It was also noted that it is not possible to automatically discover threats that haven't in the wild for a while and did not have had sufficient time to propagate. The goal in OpenCSAM is to take feedback from the user and simplify the problem.

4.4 EVENT TIMELINE EXTRACTION

Next topic in line was the question whether it is possible to extract the temporal characteristics of an event. When did a particular event happen? This question can be considered in the context of the previous assumption, that is whether we can extract from an article the date/time that an event took place? The conclusion here was that there is no way to tell when a certain article was published on the internet. One can only crawl at regular time intervals such as every couple of minutes. In some cases some articles were published in the "future" "See election result.". It is a very difficult task in general, we can only rely on the declared published date of the article and the date that the crawlers discovered the article.

4.5 TEXT SUMMARIZATION

The last topic investigated was the state-of-the-art on text summarization and how it can be improved. Can we reduce the dimensions of a text significantly while preserving the core meanings and statements? In the discussion it was mentioned that in general there are two types of text summarization: Abstractive vs extractive. Abstractive approach is the human approach as one would write an executive summary of a document, as for the extractive approach relates more to extract the most common words of phrases from the text. The dealing with this is to use existing NLP algorithms and train neural network for our corpus to show similar terms or principal terms.

4.6 CONCLUDING REMARKS

- Continue the development of OpenCSAM in annual spirals focusing on the key challenges mentioned above.
- Build a community around OpenCSAM from EU Institutions, the Academia, Member States public authorities and the private sector.

5. AREA 4- CYBER AUTONOMOUS RESPONSE, THREAT DETECTION AND SECURITY AUTOMATION

The aim of the fourth panel was to discuss developments in the field of Cyber autonomous response and threat detection in operational activities. The panel was comprised of participants working in academia and industries. More specifically, there were representatives from IBM Security (Mr. Domenico Raguseo), from S2GRUPO (Dr Luis Búrdalo & Dr. Miguel A. Juan), from University of Milan (Prof. Ernesto Damiani & Prof. Claudio Ardagna) and from Strathclyde University (Prof. Ivan Andonovic). The panel was moderated by ENISA (Dr Fabio Di Franco).

The moderator introduced the session based on the results of the previous sessions and which capability an AI system should have: an AI system should have the capability to learn, reason and think and then take action in response to what it is sensed and the planned objectives. However seems that the term AI has become a buzzword used for marketing purposes.

The terms AI might be interpreted in different way and the following schema might help in building an AI maturity model:

- Automated Intelligence: apply automation to routine processes that requires no human judgment
- Assisted Intelligence: assist people in accomplish their task providing aggregated information and automated functionality
- Augmented Intelligence: assist human judgement with a set of information which help people to make better decision (strategic and tactical/operational level)
- Autonomous Intelligence: Automating decision making processes without human intervention

As more organizations adopt a policy of continuous monitoring, security teams find themselves with voluminous quantities of monitoring and limited number of resources for the operation of the SOC. It is true that computing power and data science has progressed to the point where we can use machines to analyse the data to detect patterns and then use the patterns to create predictions and models, which are going to be tested. However, the network monitoring and detection fields have not progressed as fast as other fields (e.g. image recognition). A key point to consider is that quantity of labelled data that are available in image recognition in order to train the model is enormous and it is not available in other disciplines.

It is clear from the panel that we are still in the early stage of using an Artificial intelligence to automate decisions without human intervention in the network monitoring. Machines might assist humans in accomplishing their task providing aggregated information and automated functionality and reduce the time spent in boring and repetitive tasks.

In this respect, S2grupo shows that anomaly detection might be used in OT [4] since it is easier to train the model as the expected behaviours are more repeatable. The research they conducted

[4] OT (Operational Technology) refers to the hardware, firmware and software that either monitor or control processes and activities in the industrial sector.

in OT will be used soon to reduce the amount of low complexity alerts in order for SOC teams to focus on more high valuable tasks and to reduce the errors due to fatigue.

Another take-away from the discussion is that network environments are very different and heterogeneous. The adaptation of a machine learning model trained in one environment to another environment is not a trivial task as the University of Milan pointed out. The adaptation phases might take long: in that phase, analysts will need to be assigned to validate what is a normal behaviour or an unusual activity in that context until an Acceptable Levels of False Alarms is reached. Moreover, the evasive adversaries make the task harder: it is very hard to detect a threat and it is even harder to predict a future behaviour in order to train a model.

Examples of Machine Learning and big Data analytics used today in the detection of attacks use information from threat intelligence combined with specialized algorithms applied at different level of the kill chain[5]. In fact, machine learning algorithms might be chosen based on specific problem to solve and each stage of the kill chain (starting from the Initial exploit, through the stages of Gain persistence and Local network discovery until data exfiltration) will need to be analysed separately with specialized algorithms based on specific training data. Finding a highly specialized algorithm which might solves that specific problem is already a technical challenge based on trials and errors. Another complexity to add to the puzzle is that not all the threats will use the same attack vectors or the same tactics. Therefore, only by using a cognitive model based on the combined information from specialized algorithms and from the threat intelligence of the attack behaviours, an attack might be detected and stopped.

IBM notices that Automated Threat intelligence might be able to stop an attack before the initial access has been performed or during the execution phase. The future of Incident Analysis in Cyber Security is to automatically investigate incidents and anomalies based on insights from millions of external sources and cognitive reasoning. The role of the incident analysts will evolve and an Augmented Intelligence will provide fast and aggregated information which will help analysts to make better decision.

5.1 CONCLUDING REMARKS

- Machine Learning can provide fast and aggregated information that helps analysts to make better decision and reduce the time spent in boring and repetitive tasks.
- More research is needed for an automatic response without human interventions
- Build collaboration with research community

[5] https://attack.mitre.org/

6. CONCLUSIONS – FUTURE WORK

The event was a success given the feedback received from participants. Most of them would like to see events like this happening more often. Some representatives from Member States, like Portugal, have even volunteered to host follow up events in the future. Given this demand, it is suggested to bi-annually organise an event around the Blueprint in order to periodically evaluate progress, identify new challenges and bring together relevant communities from EU Institutions, the private sector and academia.

An important takeaway of the event has been that Artificial Intelligence and Operational collaboration in the Blueprint context are issues of particular interest by the legal world. Almost half of the audience were professionals dealing with law and its extensions into the cyber domain.

Regarding the Operational collaboration for the Blueprint, a number of EU Institutions showed interest and attended, like DG HOME, Eurocontrol and the European Council. The discussions that took place between these stakeholders revealed the need for extending the Blueprint to cover the whole of the EU Institutions and also to investigate success stories in Operational collaboration of other sectors, like the Banking and the Aviation sector in order to benefit from their experience in crisis management.

More research is required on information fusion and the produce of advice at the Operational level. The impact of large scale cyber-attacks cannot be seen as just technical. There are social, political, reputational and economic aspects as well as 2nd and 3rd order impacts that need to be accounted for. This is a topic of ongoing research and EU and ENISA should put more efforts towards this direction.

ENISA’s Open Source Cyber Situational Awareness Machine draw a lot of attention, especially from EU Institutions. It has been recognised as an innovative research and development project. Many participants expressed interest to participate to the development as beta testers and evaluators while some even expressed interest to contribute in the development process. ENISA should continue developing OpenCSAM in spirals and perhaps collaborate with researchers and the private sector for building a permanent research and development capability for Artificial Intelligence.

On the technical side, ENISA is planning to prepare a report on Artificial Intelligence and machine Learning in Operational Activities. In particular, this report will navigate on how Natural Language Processing based on OSINT and internet feeds might help in prioritizing response in a SOC and how an assisted intelligence might help to detect and monitor the network and autonomously respond to it.

XX-00-00-000-XX-X

ABOUT ENISA

The mission of the European Union Agency for Cybersecurity (ENISA) is to achieve a high common level of cybersecurity across the Union, by actively supporting Member States, Union institutions, bodies, offices and agencies in improving cybersecurity. We contribute to policy development and implementation, support capacity building and preparedness, facilitate operational cooperation at Union level, enhance the trustworthiness of ICT products, services and processes by rolling out cybersecurity certification schemes, enable knowledge sharing, research, innovation and awareness building, whilst developing cross-border communities. Our goal is to strengthen trust in the connected economy, boost resilience of the Union's infrastructure and services and keep our society cyber secure. More information about ENISA and its work can be found at www.enisa.europa.eu.

ENISA
European Union Agency for Cybersecurity

Athens Office
1 Vasilissis Sofias Str
151 24 Marousi, Attiki, Greece

Heraklion office
95 Nikolaou Plastira
700 13 Vassilika Vouton, Heraklion, Greece

enisa.europa.eu

ISBN 000-00-0000-000-0
doi: 0000.0000/000000