SOCCRATES

SOC & CSIRT Response to Attacks & Threats based on attack defence graphs Evaluation Systems

D5.2

Business Logic Modelling and Impact Analyser & Response Planner - Initial prototype

Deliverable type:	Demonstration
Contributing work packages:	WP5 Impact Quantification
Due date of deliverable:	31/08/2021
Submission date:	30/08/2021
Dissemination level:	PU

Responsible organisation:	IMT
Editor:	Christophe Kiennert (IMT)
Revision:	1.0

Abstract	This deliverable describes the initial prototype of the Business Impact Analyser, Business Logic Modelling and Response Planner components. It describes their functionalities, their integration with the SOCCRATES platform, as well as the current version of the graphical interfaces.
Keywords:	Business Logic Modelling, Business Impact Analyser, Response Planner, Containement CoA Generator

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 833481
Call H2020-SU-ICT-2018 • Innovation Action • Start date: September 1st, 2019

Management summary

The Business Impact Analyser and the Responser Planner are the two main components of the SOCCRATES platform that aim at performing impact quantification. The Business Impact Analyser performs an analysis of the operational impact of events such as attacks or mitigation deployment, while the Response Planner evaluates the financial impact of CoAs, on top of generation specific containment CoAs to handle the case of ongoing attacks, for which the ADG analyser cannot provide short term mitigations.

This deliverable describes the current state of the development of these two components, for which a full prototype is planned for March 2022. It presents the functionalities, architecture and interfaces of the two components, with illustrations of the graphical user interface. Since its development was started earlier, the Business Impact Analyser is already at a rather advanced state, while more development efforts have to be spent on the Response Planner side.

For each component, the next steps necessary to reach the implementation of the full prototype have been clearly identified, and appear to be in line with the initial planning.

Contributor(s)	Christophe Kiennert (IMT) Marwan Lazrag (IMT) Paul-Henri Mignot (IMT)

Reviewer(s)	Ewa Piatkowska (AIT)

Security Assessment	**See deliverables table for security assessment requirements**
Approval Date	17/08/2021
Remarks	No security sensitive issues.

TABLE OF CONTENTS

1 INTRODUCTION – RATIONALE OF THIS DOCUMENT 6
1.1 The SOCCRATES project 6
1.2 This deliverable 7
2 BUSINESS IMPACT ANALYSER AND BUSINESS LOGIC MODELLING 8
2.1 Component role and functionalities 8
2.1.1 Impact computation 8
2.1.2 Asset criticality evaluation 9
2.2 Component architecture 9
2.2.1 Business Impact Analyser API and Cortex analysers 10
2.2.2 Business Logic Modelling inputs 10
2.2.3 Requests to the Business Impact Analyser 12
2.3 Graphical interface of the Business Impact Analyser 14
2.3.1 Graphical interface features 14
2.3.2 Probability graph 15
2.3.3 Temporal graph 16
2.3.4 Business Impact Analyser functionalities 16
2.3.5 Graph building strategy 18
2.4 Next steps 19
3 RESPONSE PLANNER 21
3.1 Component role and functionalities 21
3.1.1 Containement CoA generation 21
3.1.2 CoA ranking 22
3.2 Component architecture 24
3.2.1 Response Planner API and Cortex analyser 24
3.2.2 Response Planner graphical user interface 26
3.3 Next steps 27
4 REFERENCES 28
5 ABBREVIATIONS 29

1 Introduction – Rationale of this document

This section introduces the SOCCRATES project and defines the goals of this deliverable.

1.1 The SOCCRATES project

SOCCRATES (SOC & CSIRT Response to Attacks & Threats based on attack defence graphs Evaluation Systems) is an EU funded project under the Horizon2020 programme that has the following main challenge:

How can SOC and CSIRT operations effectively improve their capability in detecting and managing response to complex cyber-attacks and emerging threats, in complex and continuously evolving ICT infrastructures while there is a shortage of qualified cybersecurity talent?

The main objective of SOCCRATES is to develop and implement a security automation and decision support platform (**'the SOCCRATES platform')** that will significantly improve an organisation's capability (usually implemented by a SOC and/or CSIRT) to quickly and effectively detect and respond to new cyber threats and ongoing attacks.

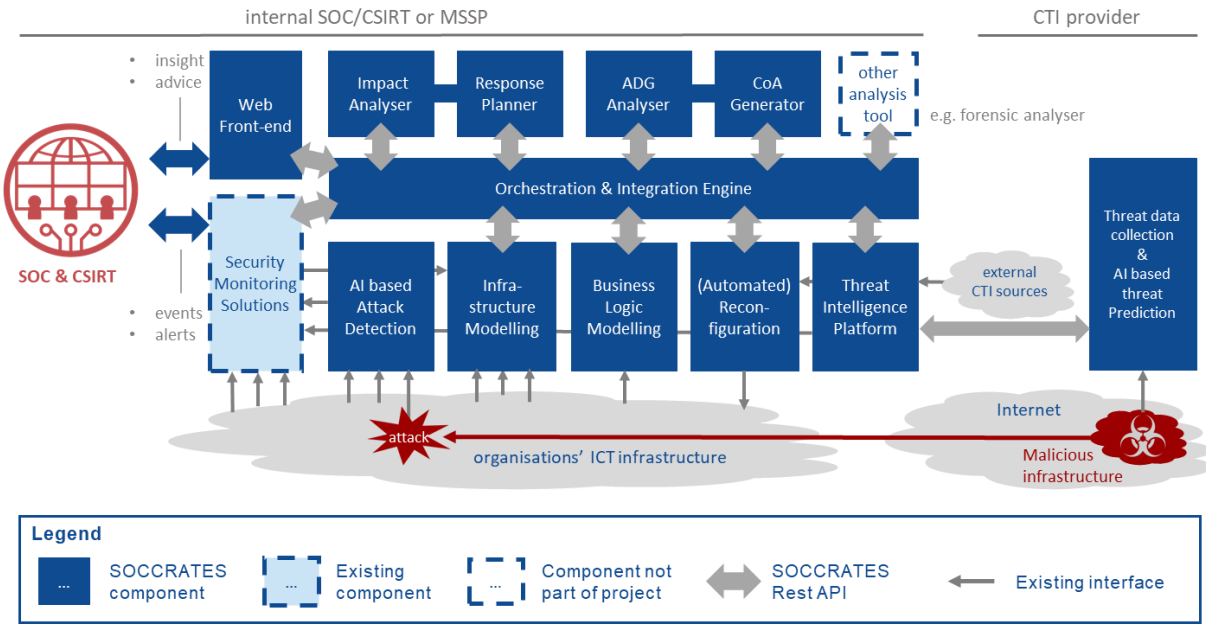

Figure 1-1 – The SOCCRATES platform

The SOCCRATES platform (see **Figure 1-1**) consists of an orchestrating function and a set of innovative components for automated infrastructure modelling, attack detection, cyber threat intelligence utilization, threat trend prediction, and automated analysis using attack defence graphs and business impact modelling to aid human analysis and decision making on response actions, and enable the execution of defensive actions at machine-speed.

SOCCRATES has the following concrete project objectives:

1. Deliver the SOCCRATES platform consisting of an orchestration function and a unique integration of innovative background solutions that seamlessly work together.
2. Show that the SOCCRATES platform can improve SOC operations by evaluating the SOCCRATES platform in two diverse real-life pilot environments.
3. Examine and illustrate the benefits of automation for selected SOC activities to help manage the cyber security skills gap in organizations.
4. Prepare for successful exploitation by the SOCCRATES partners of the individual innovated components and the integrated SOCCRATES platform in commercial products that are offered to the market and are available for the European (business) community.

Please visit www.soccrates.eu for more information on the SOCCRATES project.

1.2 This deliverable

The objective of this deliverable is to describe a first prototype of the Business Impact Analyser (and Business Logic Modelling it relies upon), as well as the Response Planner. The theoretical background of the Business Logic Modelling was already described in Deliverable D5.1 [D5.1].

The first prototype of these two components aims at providing some of the main expected functionalities, as well as a graphical interface representative of the components. For the Business Impact Analyser, the computation of the probability of impact, of the critical time, as well as the evaluation of the assets criticality have been implemented. The graphical interface aims at making it easy for an operator to identify the most critical assets, as well as the most impacted business entities when an attack targets one or several assets. However, at this stage, the component has yet to be integrated to the SOCCRATES platform, and process the inputs provided by various other components of the platform.

On the Response Planner side, the implementation was started later, in compliance with the task planning. Therefore, the component is still at an early development stage. It computes the Return on Response Investment (RORI) value of Courses of Actions (CoAs) and displays the results in a graphical interface that allows the Security Operations Center (SOC) analyst to quickly identify the most efficient CoAs. However, the evaluation of a few parameters of the RORI metric still has to be specified, using the available data from the SOCCRATES components. An interface with the Business Impact Analyser also has to be added in order to compute the business impact of CoAs. The strategy for the generation of containment CoAs has been defined, but has yet to be implemented. Finally, the output of the Response Planner will be represented as a CACAO playbook, in order to optimize the interaction with the Automated Reconfiguration component.

1.3 Structure of this deliverable

Section 2 of this document describes the initial prototype of the Business Impact Analyser, specifying the implemented functionalities, the component architecture, and the graphical interface. The section ends with a description of the next steps that will allow us to reach the full implementation of the component. Similarly, Section 0 describes the initial prototype of the Response Planner, with a structure similar to Section 2.

2 Business Impact Analyser and Business Logic Modelling

2.1 Component role and functionalities

This section describes the role and the functionalities of the Business Impact analyser component. The objective of the Business Impact Analyser is to provide the SOC analyst with an evaluation of operational impact of external events on the infrastructure, such as an attack or the deployment of a mitigation. To do so, two metrics were defined, and described in Deliverable D5.1 [D5.1]: the probability of impact on the company, and the critical time.

However, the Business Impact Analyser can also help identify the most critical assets in the infrastructure, i.e. those which contribute the most to the propagation of impacts. Therefore, on top of the evaluation of the operational impact of external events, the Business Impact Analyser can also compute a criticality value for each asset.

2.1.1 Impact computation

The impact computation is comprised of two metrics: the impact probability and the critical time.

Regarding the impact probability, the theoretical approach is based on the inclusion-exclusion principle, which is not tractable when the size of the graph grows. Therefore, the Monte-Carlo approximation, which considerably improves the computation time thanks to a tractable time complexity, was used instead. A more detailed description of the theoretical background of impact probability computation is provided in Deliverable D5.1 [D5.1].

The Monte-Carlo approximation was implemented to randomly explore the business graph a certain number of times, *ntimes,* while counting the number of times each node has been impacted by the shock events (which are the events occurring on the infrastructure, such as attacks or mitigation deployments). The shock events are the input list of perturbations whose operational impact we want to evaluate.
These are the steps of one graph exploration:

1. Initiate a queue *Q*, with the input list of Shock Events
2. Pop a node *n*, from *Q*
3. For each edge of the node *n*, try the impact probability of the edge by drawing a random number. If the drawn number is lower than the probability of impact of the tested edge, then the node at the other end of the edge is considered as impacted. If a node is impacted and was never added to the queue *Q*, then add it to *Q*, and increase the impact counter for this new node
4. While the queue *Q* is not empty, repeat from step 2

Eventually the approximated probability of business impact is computed by taking the mean of the impact counter (Impact counter divided by *ntimes*).

The precision of the approximated probability of impact depends on the number of graph exploration (*ntimes*). For the project, *ntimes* was set to 1000, giving an approximation with a $\pm 3\%$ precision.

The Monte Carlo algorithm was implemented in Python using standard libraries, among with multiprocessing libraries to speed up the computation of the asset criticality.

The API of the Business Impact Analyzer is coded in Python, using the external libraries Numpy, Uvicorn[1] and FastAPI[2].

Regarding the critical time, since it is equivalent to a shortest path in the mission dependency graph, the Dijkstra algorithm has been implemented.

2.1.2 Asset criticality evaluation

The Business Impact Analyser component will also focus on the assets criticality, in order to identify the assets that contribute the most to the propagation of impacts on the business functions and processes.

The criticality of an asset was estimated by computing the value of the impact probability on the business company (BC) node when a shock event occurs on this asset with a local conditional probability equal to 1. Therefore, the criticality of the assets is evaluated as a number between 0 and 1.

In order to provide a convenient visualization of the assets criticality, three levels of criticality have been defined: high, medium and low, based on the computed value. Details on the graphical interface of the component are provided in Section 2.3.

2.2 Component architecture

The internal architecture of the Business Impact Analyser is represented in Figure 2-1. A detailed description of the various elements of this architecture was provided in Deliverable D2.2 [D2.2].

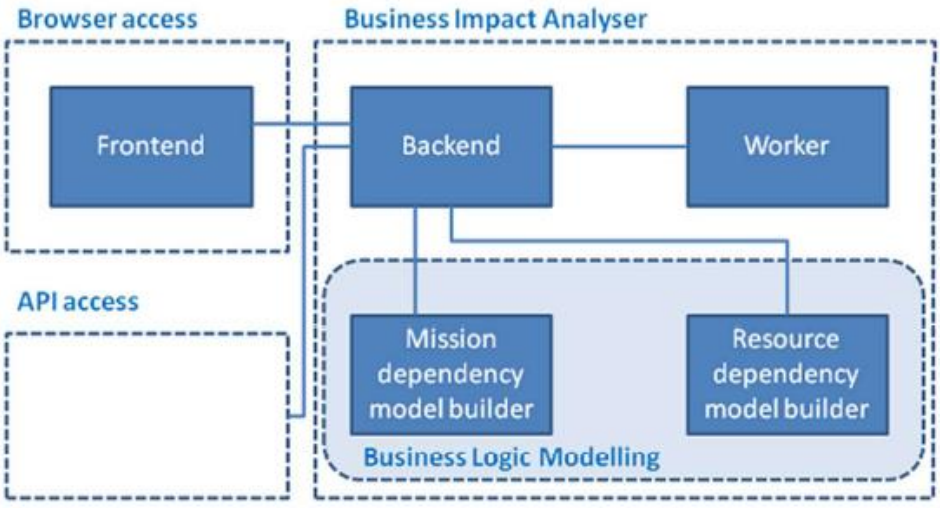

Figure 2-1 – Architecture of the Business Impact Analyser

[1] https://www.uvicorn.org
[2] https://github.com/tiangolo/fastapi

Before performing any computation, the Business Impact Analyser needs to build a model. To do so, it relies on the Business Logic Modelling, which builds a mission dependency model and a resource dependency model. These two models are then fused into one graph that represents the current model.

2.2.1 Business Impact Analyser API and Cortex analysers

The Business Impact Analyser backend runs with a Python REST API. Two Cortex analysers and one Cortex responder have been developed for the integration in the SOCCRATES platform:

- One Cortex analyser to compute the business impact of an external event
- One Cortex analyser to retrieve the critical assets (assets with a criticality value labeled as high)
- One Cortex responser to update the current model after the addition of a new system, specific to Use Case 5.

The details of the Cortex analysers and responder are provided in Deliverable D2.3 [D2.3]. Figure 2-2 displays the interactions of the Business Impact Analyser with the orchestrator (OIE).

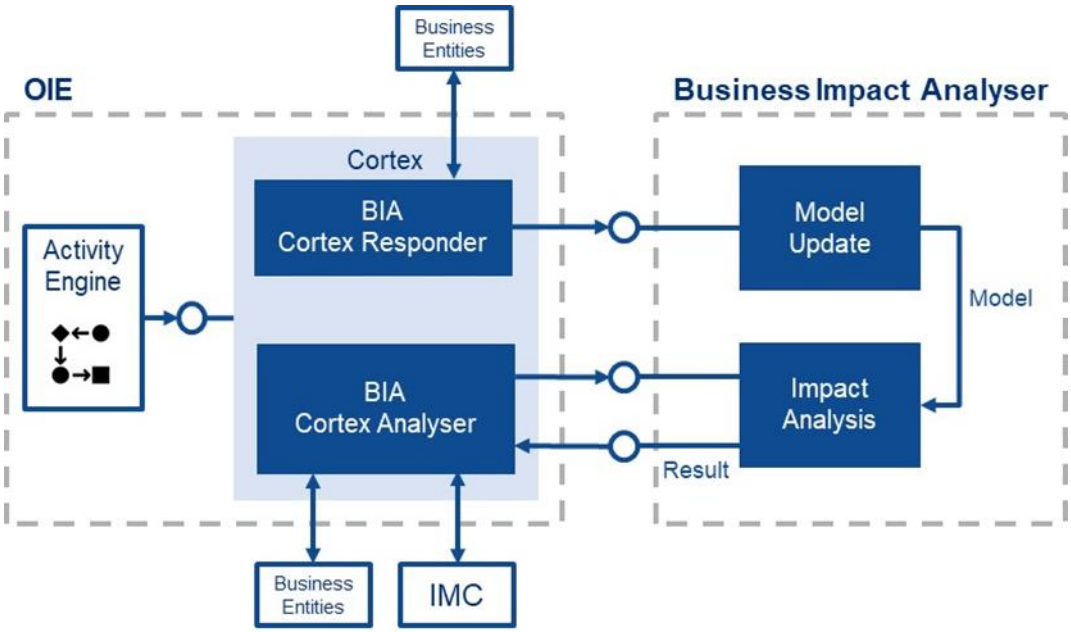

Figure 2-2 – Business Impact Analyser Cortex interfaces

2.2.2 Business Logic Modelling inputs

This sub-section details the inputs needed to initialize and update a business logic model. As stated previously, and as detailed in Deliverable D5.1 [D5.1], a model is comprised of two parts :

- The mission dependency model, which contains the graph of the business functions and processes interdependencies
- The resource dependency model, which contains the graph of asset interdependencies

2.2.2.1 Mission Dependency Model

To build the mission dependency model, the necessary input is a graph representing the company's business functions and processes along with their interdependencies. It is also necessary for this graph to feature the assets that directly support business functions. An example of such graph is provided in Figure 2-3.

In the context of the SOCCRATES project, this graph has to be manually built, as there is no easy method to automate this, and developing one would require substantial work that would be out of the project scope. However, as advocated in Deliverable D5.1 [D5.1] this graph is not likely to change over time, or only with minor tweaks, since the business organization of a company is much more stable than its infrastructure. Therefore, even though a lot of effort is needed on the company side to build of first version of the graph for the initialization of the component, the model updates will remain very light.

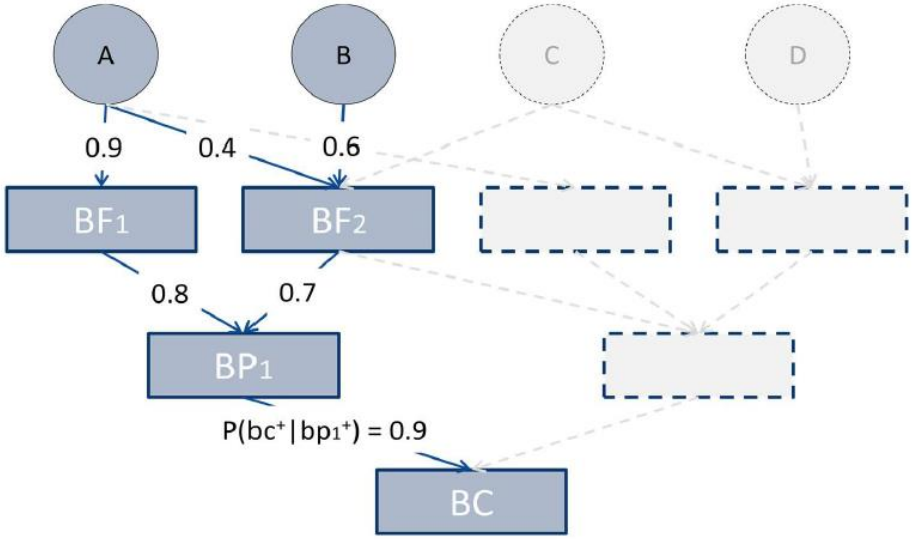

Figure 2-3 – Mission dependency model

To initialize the mission dependency model, the Business Logic Modelling requires such a graph, represented as a matrix describing the interdependency values, such as the one represented in Table 2-1 which contains the impact probabilities. Another matrix is required to provide the downtime tolerances between business entities.

As specified in Deliverable D5.1[D5.1], this graph has a few constraints, such as the need of a unique Business Company (BC) node acting as a sink for the graph. It must also be acyclic, in accordance with a tree structure. For the initialization phase, the matrix will be provided to the component beforehand, as a CSV file.

Node	BF1	BF2	BP1	BC
A	0.9	0.4	0.0	0.0
B	0.0	0.6	0.0	0.0
BF1	0.0	0.0	0.8	0.0
BF2	0.0	0.0	0.7	0.0
BP1	0.0	0.0	0.0	0.9
BC	0.0	0.0	0.0	0.0

Table 2-1 – Business Entity probability matrix

In the context of the SOCCRATES project, the mission dependency model may only be updated during Use Case 5, where a new system is added to the infrastructure. However, if the newly added system does not bring any change to the business entities, then only the resource dependency model will be updated.

2.2.2.2 Resource Dependency Model

The notion of dependency between assets is defined in relation with the quantity of data exchanged between assets, as described in more details in Deliverable D5.1 [D5.1]. The amount of exchanged data will be retrieved every 24 hours with a request to the Infrastructure Modeling Component (IMC) , where Netflow data will be gathered.

The traffic matrix deduced from the Netflow data will be transformed into a probability matrix, from which the graph can be deduced. As of now, the Netflow adapters are not yet ready to be used, therefore a local CSV file was created to represent the traffic matrix.

Table 2-2 presents an example of probability matrix deduced from a traffic matrix, in such a way that the sum on each column is equal to 1.

Node	A	B	C	D
A	0.0	0.6	0.1	0.2
B	0.1	0.0	0.0	0.3
C	0.2	0.0	0.0	0.5
D	0.7	0.4	0.9	0.0

Table 2-2 – Inter-assets probability matrix

Once a mission dependency model and a resource dependency model are built, the matrixes are fused into a single adjacency matrix, which corresponds to the current business logic model. Table 2-3 presents the adjacency matrix built from the matrixes from Table 2-1 and Table 2-2.

Node	A	B	C	D	BF1	BF2	BP1	BC
A	0.0	0.6	0.1	0.2	0.9	0.4	0.0	0.0
B	0.1	0.0	0.0	0.3	0.0	0.6	0.0	0.0
C	0.2	0.0	0.0	0.5	0.0	0.0	0.0	0.0
D	0.7	0.4	0.9	0.0	0.0	0.0	0.0	0.0
BF1	0.0	0.0	0.0	0.0	0.0	0.0	0.8	0.0
BF2	0.0	0.0	0.0	0.0	0.0	0.0	0.7	0.0
BP1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.9
BC	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 2-3 – Adjacency matrix with conditional impact probabilities

Since there is no downtime tolerance for inter-asset impacts (the impact is supposed to propagate immediately since it does not depend on the business organization), there is no downtime tolerance matrix expected as input for the resource dependency model.

2.2.3 Requests to the Business Impact Analyser

After building a current model, comprised of a Mission Dependency model and a Resource Dependency model, the Business Impact Analyser can receive requests to:

- compute the business impact of an external event
- evaluate the asset criticality

The first type of request requires the following inputs:

- the node(s) impacted by the event
- the level of compromise in confidentiality, integrity and availability, evaluated as real numbers between 0 and 1

Figure 2-4 displays a shock event impacting node B of the current model, with a probability of 0.3 deduced from the level of compromise. The computation for the business impact will therefore trigger the calculation of the impact probability on the BC node. Similarly, a shortest path algorithm will be applied on the downtime tolerance graph to provide the critical time.

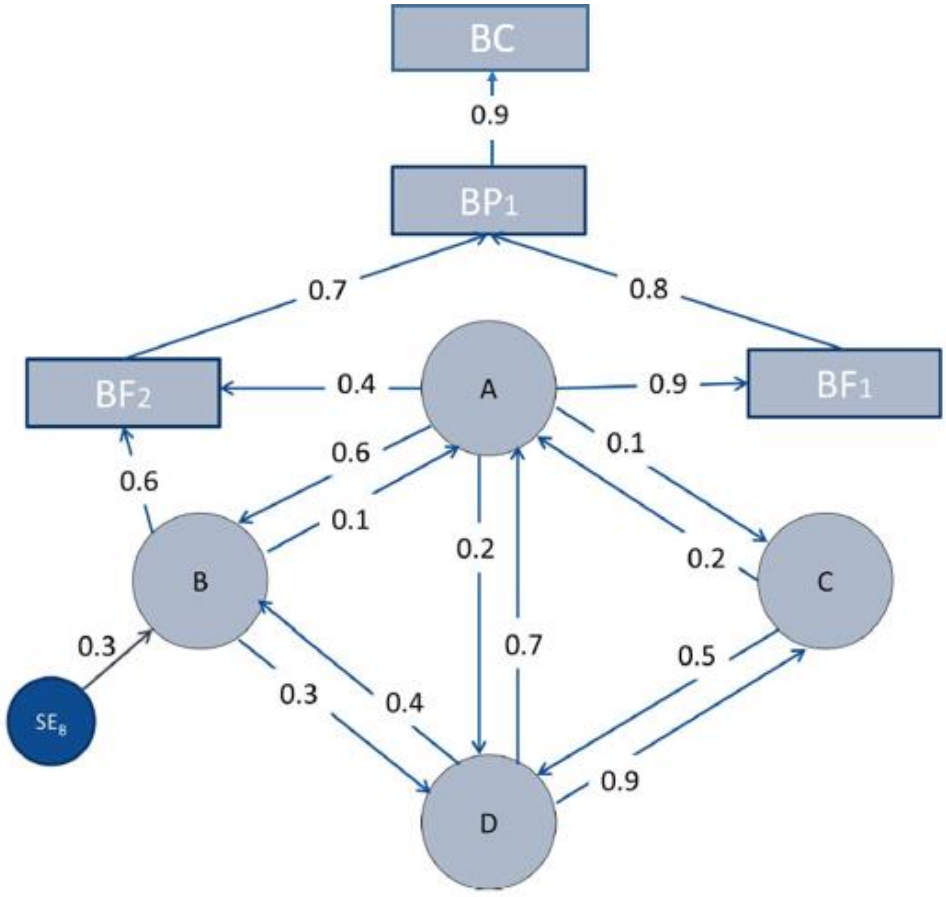

Figure 2-4 – Business logic model with conditional impact probabilities

The second type of request does not require any additional input. To evaluate the criticality of a node, an impact probability request is sent while applying a shock event on that node with an impact probability equal to 1.

2.3 Graphical interface of the Business Impact Analyser

This section describes the current status of the Business Impact Analyser component with its graphical interface. In particular, it displays the various functionalities of the component:

- The probability graph
- The temporal graph
- The assets criticality
- The impact probability of a shock event on the business functions and processes

2.3.1 Graphical interface features

The Business Impact Analyzer Frontend is a Web Graphical User interface implemented in Javascript, running on a Python web server. Illustrations of this interface are presented in Sections 2.3.2 onwards.

The graphical interface of the Business Impact Analyser is running on a server to provide two functionalities :

- The calculation of the impact on the business function and process : impact probability and critical time
- The computation of the criticality of assets

Several buttons have been developed in the graphical interface to help the SOC analyst to assess the situation and identify the most vital assets and business processes:

- Graph type: probability graph or temporal graph
- Get impact button: compute the impact probability on the business functions and processes
- Get criticality button: compute the assets criticality

2.3.1.1 Interactions with the graph

A node can be added with *Shift + Click* on an empty space. Then it is possible to select the type of this new node between Asset, Business Function and Shock Event.

An edge between two nodes can be added by *Shift + Click* on a first node then hold and drag to a second node. The direction of the edge created is from the first node to the second one. Then it is possible to enter the forward and backward impact probability between the first node and the second node.

A node or an edge can be edited by *Ctrl + Click* on it. A node or an edge can be deleted by selecting it and using *Delete* key.

It is possible to import or export a graph in JSON format by clicking on the corresponding icons at the bottom of the screen.

2.3.1.2 Sources

The implementation of the Business Impact Analyser (BIA) Frontend relies on several resources, coded in JavaScript.

It is mainly based on the *Mindcraft* code from edipa [https://github.com/eldipa/Mindcraft].

It also makes use of the following projects:

- d3: [https://d3js.org/d3.v3.min.js], Version: 3.5.17, BSD 3-clause "New" or "Revised" License
- FileSaver: http://purl.eligrey.com/github/FileSaver.js], Version: 1.3.2, MIT License
- Javascript-Undo-Manager: [https://github.com/ArthurClemens/Javascript-Undo-Manager], Version: 1.0.5., MIT License
- JavaScript Canvas to Blob: [https://github.com/blueimp/JavaScript-Canvas-to-Blob/blob/master/js/canvas-to-blob.js],Version: 3.2.0, License: MIT License
- AlertifyJS: [https://github.com/MohammadYounes/AlertifyJS], Version: 1.11.0,GNU GENERAL PUBLIC LICENSE Version 3
- Superformula Tweening: [https://bl.ocks.org/mbostock/1020902], GNU GENERAL PUBLIC LICENSE Version 3

2.3.2 Probability graph

Figure 2-5 displays the assets, the business functions and processes and the business company, as well as the impact probability dependencies between these nodes. The represented graph is the same as the example provided in Figure 2-4.

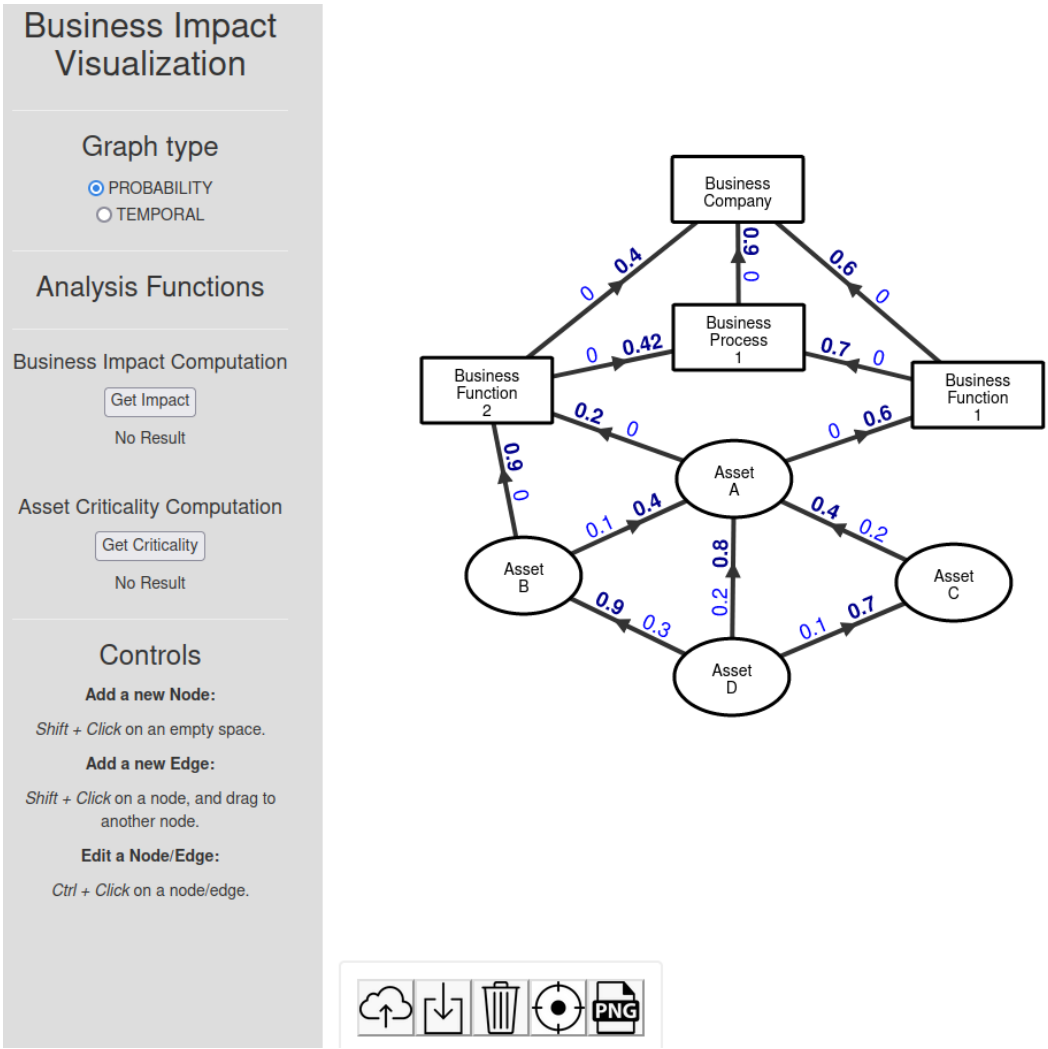

Figure 2-5 – Probability graph

In order to lighten the visual density of the graph, bidirectionnal impacts are represented with a single directed arc featuring two values: one colored in dark blue, and one colored in light blue. The dark blue value corresponds to the direction featured by the arrow on the arc, while the light blue value corresponds to the backward direction. For example, in Figure 2-5, the impact probability from node D to node B is 0.9, while the impact probability from node B to node D is 0.3. In accordance with the constraints specific to the mission dependency model, this graph contains no backward dependency on the arcs between business entities.

2.3.3 Temporal graph

Figure 2-6 presents the same graph but displaying the downtime tolerances instead of the conditional probabilities. The representation rules of the dependencies described in the previous section apply in this graph as well.

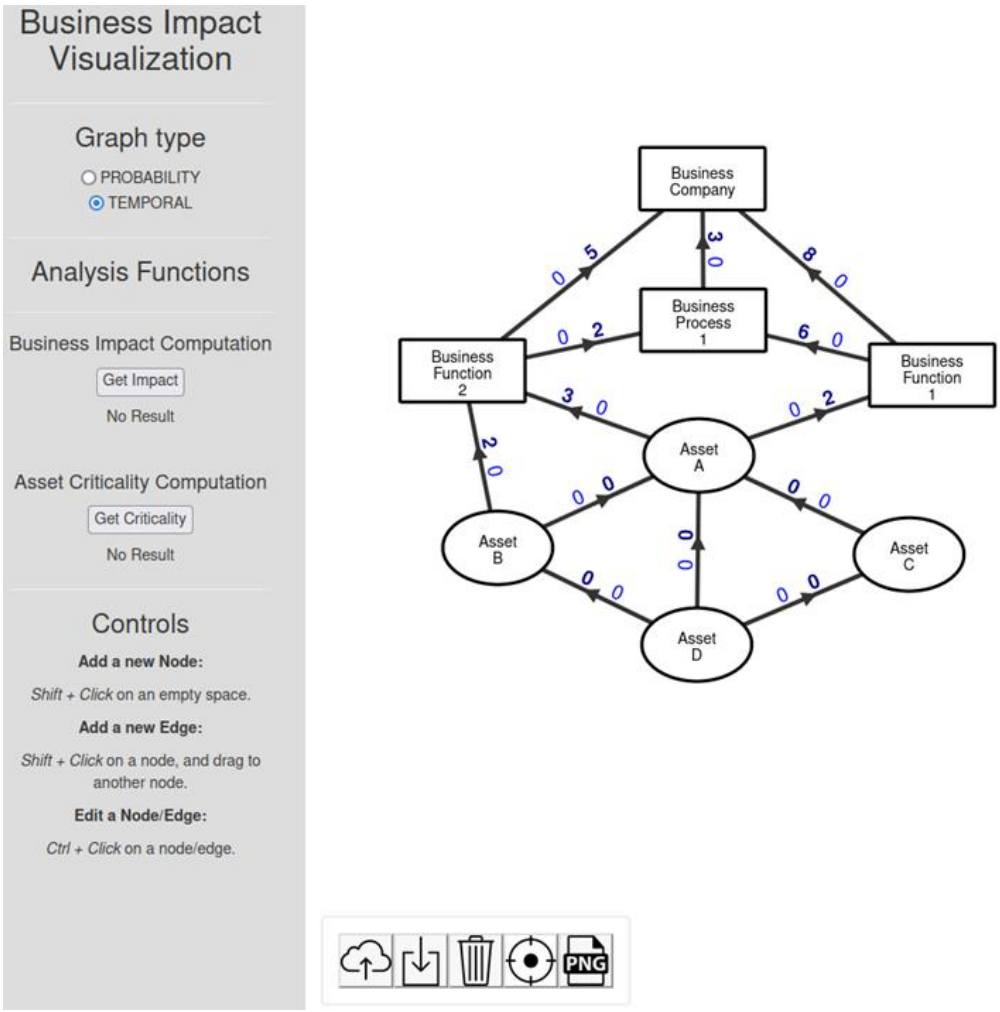

Figure 2-6 – Temporal graph

2.3.4 Business Impact Analyser functionalities

2.3.4.1 Asset criticality computation

The BIA is able to calculate the criticality of assets, which returns the assets that contribute the most to propagate the impact on the most important business function (business company node).

The computed criticality is displayed in orange above each Asset node. Then, depending on the value of the impact, the node is colored:

- in green, if the value is between 0 and 1/3, because it is considered as a low critical asset
- in orange, if the value is between 1/3 and 2/3, it is considered as a medium critical asset
- in red, if the value is between 2/3 and 1, it is considered as a highly critical asset.

The value ranges for coloring can be freely configured.

Figure 2-7 displays a screenshot of the graphical interface after the computation of the asset criticality values.

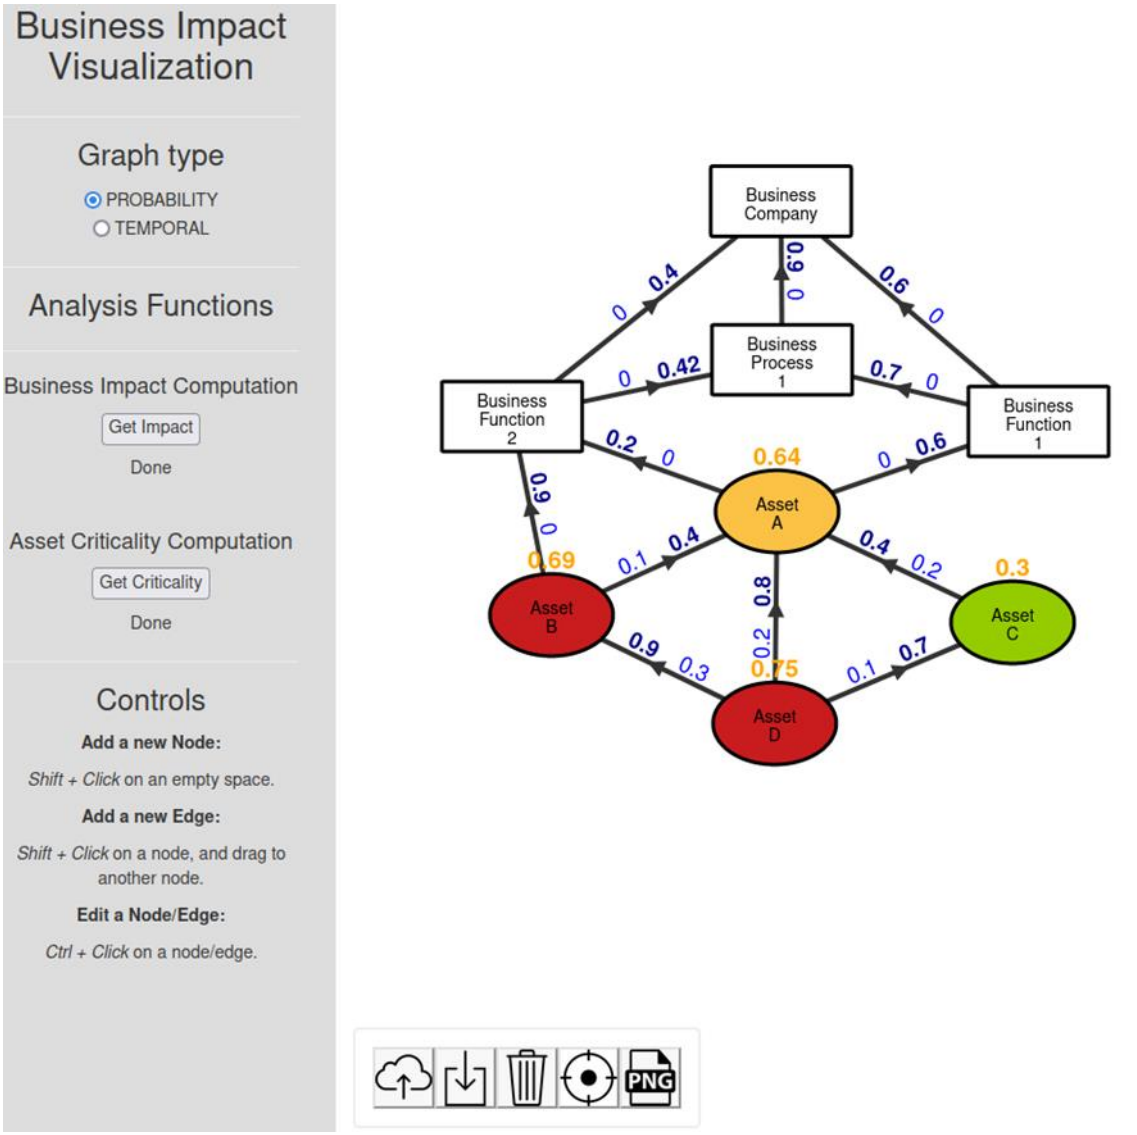

Figure 2-7 – Asset criticality computation

2.3.4.2 Impact probability computation

We have implemented the Monte-Carlo approximation to compute the impact probability of a shock event on the business functions and processes.

The computed business impact is displayed in red above each Business Function node. Then, depending on the value of the impact, the node will be colored following the same rule as the one used for the asset criticality computation.

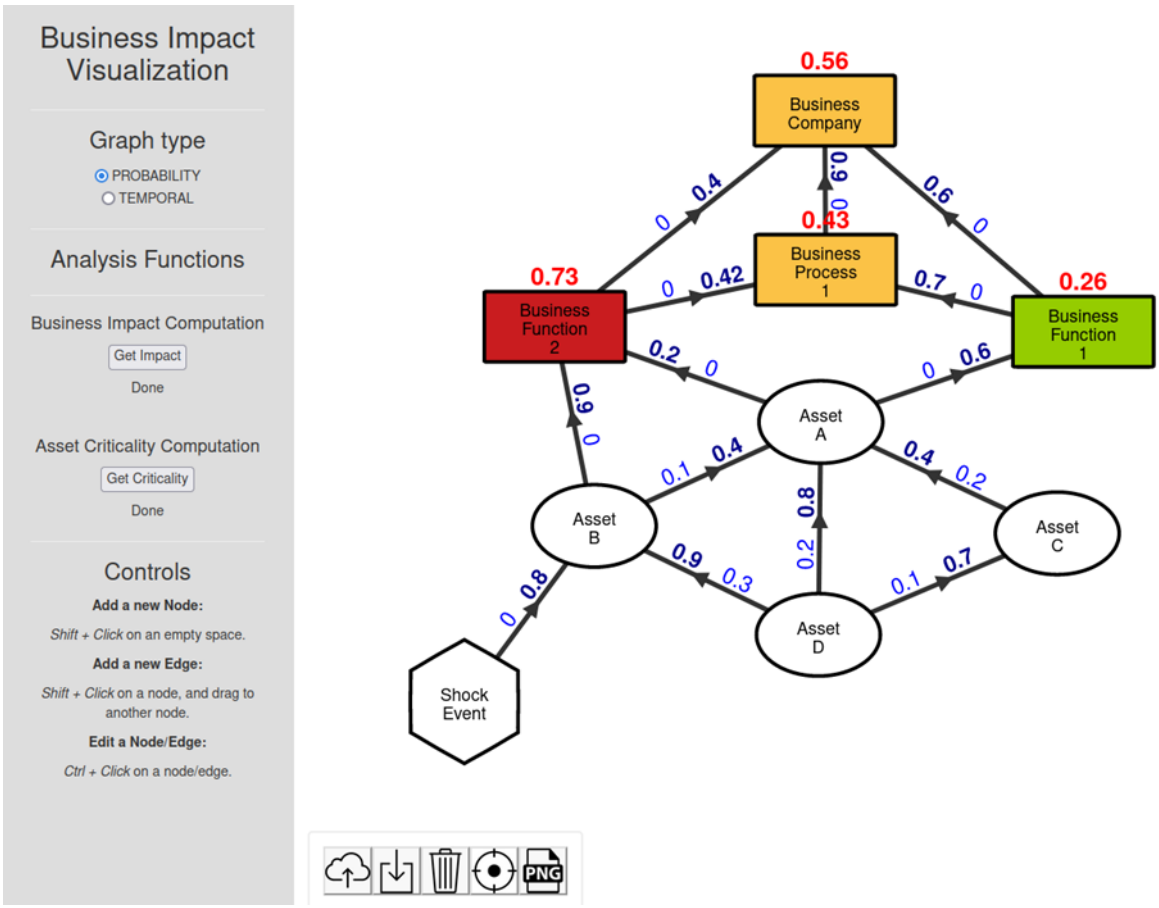

Figure 2-8 – Impact probability computation

In the example of Figure 2-8, the shock event is presented as follows : (Shock Event, Asset B, 0.8). Three levels of impact of this shock event can be observed on the business functions and processes and the business company. A high impact on Business Function 2, a moderate impact on Business Process 1 and Business Company, and a low impact on Business Function 1.

2.3.5 Graph building strategy

In the graphical user interface, a graph can be built from the adjacency matrix constructed beforehand by the BIA backend. However, the generated graph can be quite difficult to read in the case of large

graphs. Although the user interface makes it possible to drag and drop the nodes to improve readability, it was deemed necessary to think of a strategy so that the generated graph naturally highlights the important information, i.e. the business entities.

Therefore, the BIA displays the nodes with a hierarchy, starting with the Business Company node on the top, and then successively display business function nodes and assets in several circular arcs.
When adding an edge between two nodes from the adjacency matrix, the direction of the edge is chosen according to the highest impact probabilty between the two nodes, in case there is bidirectional impact. For example, considering two assets A and B with the probability of impact from A to B being 0.2 and the probability of impact from B to A being 0.8, then the edge will be drawn with an arrow from B to A, with the value 0.8 in bold dark blue, and the value 0.2 in light blue above the edge.

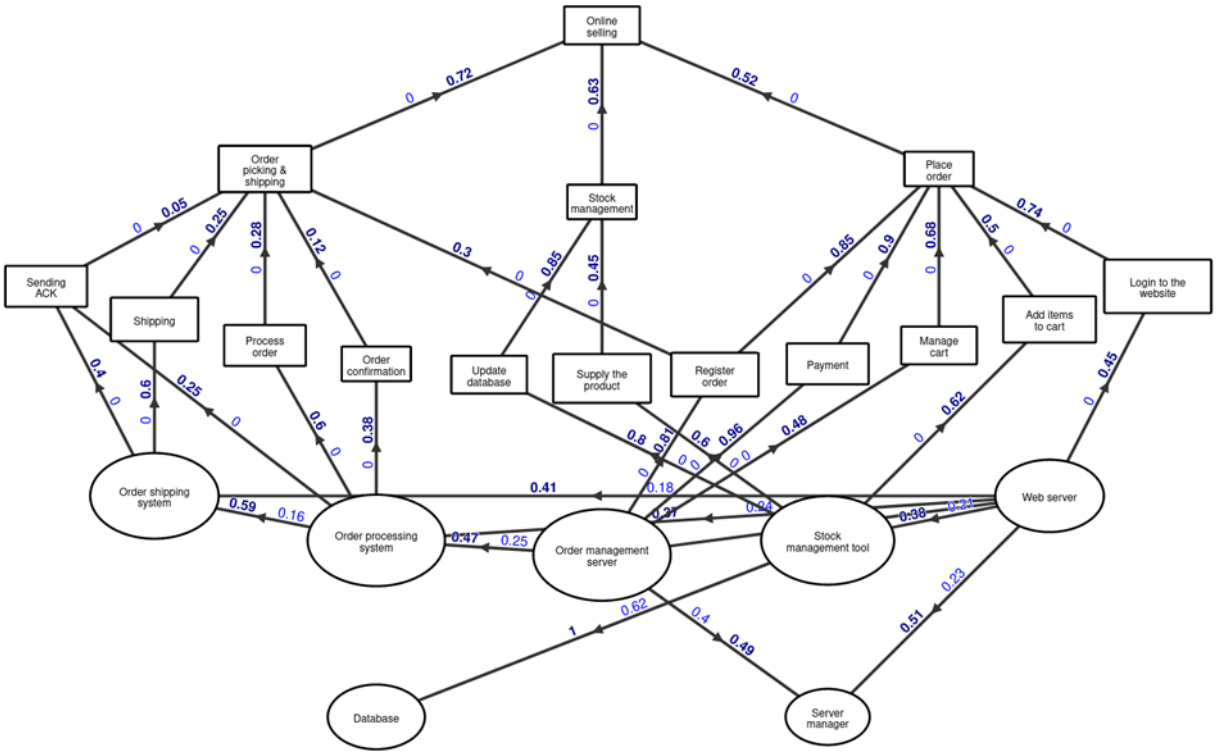

Figure 2-9 – Graph generated from an adjacency matrix

The graph must have a Business Company node, which is a Business entity with no outgoing edge. In other words, it is impacted by other nodes, but does not impact any node. Without a Business Company node, the GUI will not be able to compute the criticality values.

Figure 2-9 displays an example of generated graph. As shown in this figure, the hierarchical structure of the business functions, business processes and business company can be easily understood even without any manual interaction with the nodes.

2.4 Next steps

The development of the Business Impact Analyser is currently at an advanced stage. The next steps will focus on the finalization of the graphical interface and on the integration into the SOCCRATES platform. More precisely, the next steps will include:

- processing of the Netflow data to build a traffic matrix, from which the resource dependency model can be deduced
- addition of shortest path visualization in the graphical interface upon computing the critical time on the temporal graph
- definition of the mission dependency graphs with the pilot sites
- ensuring the scalability of the computations

Regarding the last item, substantial work has already been done with respect to code optimization, leading to very satisfying computation times on reasonably large graphs. Depending on the size of the graphs generated from the input data of the pilot sites, further optimization might be necessary, but the current performances are promising.

3 Response Planner

3.1 Component role and functionalities

In the context of Use Case 1: Response on Detected Ongoing Attack, the Response Planner (RP) is a component that will suggest a list of containment strategies, i.e. Courses of Actions (CoAs) to mitigate the current threat.

In all the use cases, following up CoA generation, the Response Planner (RP) will propose to a human analyst a ranked list of CoAs with metrics based on financial and operational impact values, in order to allow the analyst to choose the optimal CoA.

Thus, the RP provides two main functionalities: containment CoA generation and CoA ranking.

3.1.1 Containement CoA generation

The scenario of use case UC1 describes an ongoing attack that requires a quick response. However, due to the lack of threat information, the Attack Defense Graph (ADG) component will not be able to provide valuable analysis for a short term response, thus leaving the CoA Generator unable to generate CoAs dedicated to the containment of the attack. In order to mitigate this ongoing attack, the Response Planner will generate specific CoAs, called Containment CoAs. These Containment CoAs are generated from threat information about the ongoing threat obtained from the detection modules, and aims to be simple and short term countermeasures.

3.1.1.1 Strategy for containment CoA generation

The actions of a Containment CoA can be either host isolation or blocking traffic. To generate Containment CoAs, the RP receives as input a SIEM event, one or more source IP addresses and one or more destination IP addresses. The SIEM event will be given by the REASENS module or by the pilot site SIEM. The range of SIEM events has been restrained to three labels that the RP will be able to process.

The first label is *Initial Access.* This SIEM event means that an attacker managed to access one or more hosts, which will be labeled as compromised. The event will include the IP address of the compromised host. Depending on the attack technique detected, the compromised host could be given as a source or destination IP address. The RP will treat an internal IP address, in the source or destination IP address field, as a compromised host, and will generate a Containment CoA to isolate it.

The second label is *Lateral Movement.* This SIEM event alerts that suspicious communications between a compromised host and other hosts have been detected as lateral movement of the attacker, meaning that the attack is spreading to other hosts. The event will include one or more source IP addresses and one or more destination IP addresses. In the context of Containment CoAs, we assume that all hosts with reported sources and destination IP addresses are compromised, and should be contained. The generated Containment CoA will isolate the list of compromised host.

The third label is *Exfiltration.* This event means that one internal host is exfiltrating data outside of the company. The event will include one source IP address and one or more external destination IP addresses. The generated Containment CoAs will aim to block communication from the source IP address to destination IP address(es) and to isolate host with the source IP address.

3.1.1.2 Isolation perimeter

Depending on the situation, the analyst may prefer to deploy different containment strategies. In that way, the RP should produce several Containment CoAs corresponding to these strategies. Two strategies have been established based on a perimeter of isolation. The first one, the default one, is to isolate all compromised hosts. The other one is more preventive and aims to isolate all compromised hosts plus neighbouring hosts, acting on a bigger perimeter. These containment perimeter strategies are still being discussed.

3.1.2 CoA ranking

The second functionality of the Response Planner is to generate a ranking of the CoAs produced by the CoA Generator, in order to select the most suitable CoA to apply. The RP will perform an analysis of the financial and operational impact of each CoA.

3.1.2.1 Financial impact

The financial impact of a CoA is quantified by the Return on Response Investment (RORI) metric. The *RORI* metric is based on four parameters: Annual Loss Expectancy (ALE), Risk Mitigation (RM), Annual Infrastructure Value (AIV), Annual Response Cost (ARC), and computed as follows :

$$RORI = \frac{ALE \times RM - ARC}{ARC + AIV} \times 100$$

The Risk Mitigation (RM) parameter is the expected percentage of the mitigation of the attack by the CoA. In the context of the SOCCRATES project, the estimation of this parameter is based on Time to Compromise (TTC) values, computed by the Attack Defense Graph. A set of key assets will be used as a base to compute an initial TTC ($initialTTC$), and the same key assets will be used to compute the expected new TTC ($newTTC$) after the CoA is applied. These two TTC values will be provided to the RP by the CoA Generator.
Based on the initial and new TTC values, the Risk Mitigation is estimated as follows:

$$RM = 1 - \frac{initialTTC}{newTTC} = \frac{\delta TTC}{newTTC}$$, knowing that $newTTC \geq initialTTC \geq 0$

If $newTTC < initialTTC$, it means that the CoA is making the infrastructure more vulnerable, so it should not be applied.
We can observe that if the CoA succeeds to double the $initialTTC$ ($newTTC = 2 \times initialTTC$), then the RM will be equal to 50%.

The Annual Loss Expectancy parameter is the expected loss from the attack if no CoA is deployed. The estimation of this parameter remains to be defined, but is likely to be defined as a function of the type of SIEM event, of the criticality value of the targeted asset, and possibly of the $initialTTC$ value.
The Annual Response Cost parameter is the monetary deployment cost of the CoA, and will be provided by the CoA Generator.

The Annual Infrastructure Value parameter will be provided beforehand by the pilot site during the initialization phase, by evaluating the monetary value of its infrastructure. This value should be evaluated once, and will then remain stable. In real world deployment, it might be updated periodically, but since it mainly acts as a scaling parameter in order to adapt the computed RORI values to the size of the company, it has no real impact on the ranking of the CoAs.

In use case UC1, because of the limited knowledge on the ongoing threat, Containment CoAs will be treated differently by the RP. More specifically, only the operational impact will be computed, since the RORI value is not relevant for these. Indeed, the computation of financial impact is really meaningful for long term mitigations only, as short term mitigations will not induce important investment or maintenance costs.

3.1.2.2 Operational impact

The operational impact of a CoA is estimated with the help of the Business Impact Analyser component. The objective is to evaluate how much the deployment of a CoA will impact the business company.
The details of the request to the Business Impact Analyser and the computation of the business impact remain to be discussed.

3.1.2.3 CoA ranking output

The output of the CoA Ranking function will be the list of evaluated CoAs ranked by according to either the RORI value or the operational impact value. The defenses of the CoAs will be represented as a CACAO playbook [CACAO] along with a human readable label.

The RP will automatically transform the defenses list of the input CoA into a CACAO playbook, containing OpenC2 commands [OpenC2] or human readable commands. Using CACAO playbook allows us to standardize the description of the CoA defenses, while being able to mix automatic OpenC2 commands and manual human commands, in anticipation of the upcoming interactions with the Automated Reconfiguration component.

An example of CACAO playbook is provided in Figure 3-1.

```
{"type": "playbook",
 "id": "playbook--soccrates-coa-123456123456",
 "name": "Soccrates CoA Playbook",
 "description": "This playbook will apply the CoA in the system",
 "playbook_types": [
   "prevention"
 ],
 "created_by": "identity--uuid",
 "created": "2021-07-19T23:32:24.399Z",
 "labels": [
  "Patch",
  "2FA",
  "CoA"
 ],
 "features": {
  "parallel_processing": true,
 },
 "workflow_start": "start--soccrates-coa-123456123456",
 "workflow": {
  "start--soccrates-coa-123456123456": {
   "type": "start",
   "on_completion": "parallel--soccrates-coa-123456123456"
  },
  "parallel-soccrates-coa-123456123456": {
   "type": "parallel",
   "name": "Apply defenses of the Coa",
   "next_steps": [
    "defense1--soccrates-coa-123456123456",
    "defense2--soccrates-coa-123456123456"
   ]
  },
  "defense1--soccrates-coa-123456123456": {
   "type": "single",
   "name": "Patch Application",
   "description": "This step will Patch the application on selected hosts",
   "on_completion": "end--soccrates-coa-123456123456",
   "target": "112233",
   "commands": [
    {
     "type": "manual",
     "command": "Download and install the patch for the application in selected hosts"
    }
   ]
  },
  "defense2--soccrates-coa-123456123456": {
   "type": "single",
   "name": "Two Factor Authentication",
   "description": "This step will enable Two Factor Authentication",
   "on_completion": "end--soccrates-coa-123456123456",
   "target": "445566",
   "commands": [
    {
     "type": "manual",
     "command": "Enable Two Factor Authentication on the selected hosts"
    }
   ]
  },
  "end--soccrates-coa-123456123456": {
   "type": "end"
  }
 }
}
```


Figure 3-1 – Example of a CACAO playbook for a CoA

3.2 Component architecture

3.2.1 Response Planner API and Cortex analyser

The Response Planner will run with a Python REST API. This API has been developed using the external libraries Uvicorn[3] and FastAPI[4].

The Response Planner has two Cortex Analyzers, one for each of its functions:

- SOCCRATES_RP_CONTAINMENT_COA_GENERATOR
- SOCCRATES_RP_COA_RANKING_ANALYZER

These Cortex Analysers have been developed using the external libraries *CortexUtils*[5] and *Requests*[6]. The first one aims at generating containment CoAs based on attack detection data provided by REASENS, which is part of the AI-based Anomaly Detection (AAD) or by the pilot site SIEM. The second Cortex Analyser aims at computing the financial and operational impacts of CoAs in order to obtain a CoA ranking that will help that CoA operator in decision making.
The details of these analysers are provided in Deliverable D2.3 [D2.3].

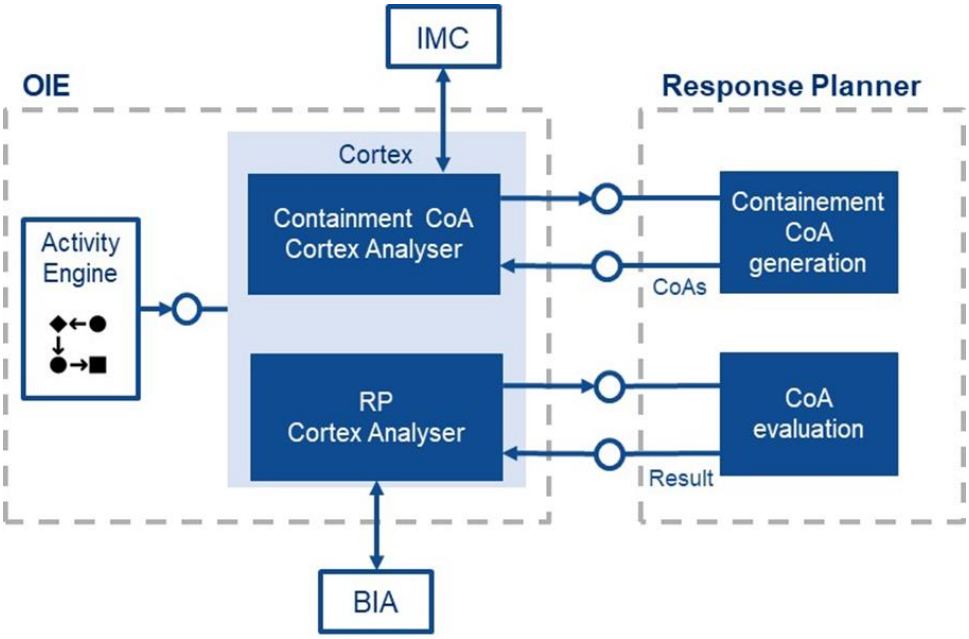

Figure 3-2 – Response Planner API interaction with Cortex Analyzers and other components

The orchestrator (OIE) interacts with the Response Planner API through the Response Planner Cortex Analysers, as shown in Figure 3-2.

The Response Planner will interact with the Business Impact Analyzer, to compute the operational impact of a CoA.

The Response Planner may interact with the IMC to obtain a precise list of security controls, in order to provide specific OpenC2 commands in the CACAO playbooks.

3 https://www.uvicorn.org/
4 https://github.com/tiangolo/fastapi
5 https://cert-bund-cortex-analyzers.readthedocs.io/en/latest/cortexutils3.html
6 https://docs.python-requests.org/en/master/

3.2.2 Response Planner graphical user interface

The Response Planner has a JavaScript Web Graphical User Interface that summarizes all the parameters relevant to establish the CoA ranking. These parameters can be manually edited in order to allow a dynamic RORI recomputation if needed.

There are three input fields outside of the table, which contain parameters that do not depend on a specific CoA : The Annual Infrastructure Value (AIV), the Annual Loss Expectancy (ALE) and the Initial Time to Compromise. It is possible to edit any of these values and update the *RORI* calculation by clicking on the *Update* button. Figure 3-3 presents the interface with the compact view of the CoAs (as opposed to the full display presented in Figure 3-4).

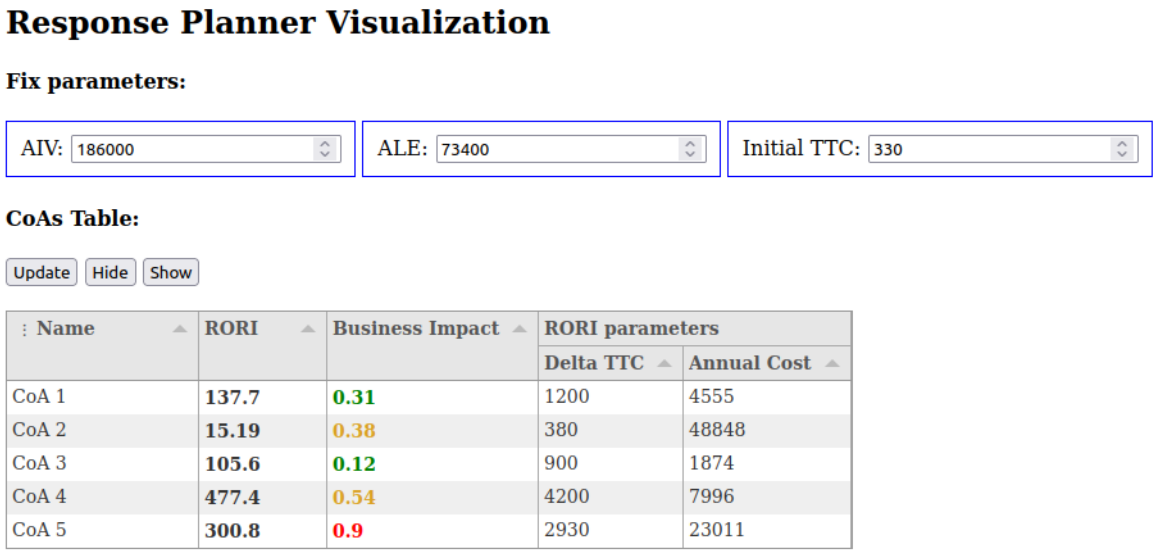

Name	RORI	Business Impact	RORI parameters	
			Delta TTC	Annual Cost
CoA 1	137.7	0.31	1200	4555
CoA 2	15.19	0.38	380	48848
CoA 3	105.6	0.12	900	1874
CoA 4	477.4	0.54	4200	7996
CoA 5	300.8	0.9	2930	23011

Figure 3-3 – Response Planner Graphical User Interface, compact CoA display

The Business Impact field is the value returned by the BIA when evaluating the CoA business impact on the company.
The values are colored following the same coloring rule applied in the BIA (in Section 2.3.4.1):

- green if the value is between 0 and 1/3, considering that the CoA has a low operational impact
- orange if the value is between 1/3 and 2/3, meaning a medium operational impact
- red, if the value is between 2/3 and 1, meaning a high operational impact

The Delta TTC field is the difference between the new TTC, expected after the deployment of the CoA, and the Initial TTC. The delta should be positive since the CoA has to reinforce the security of the infrastructure and slow down the attack.
The Annual Cost field is extracted from the cost of the CoA provided by the CoA generator, and has to be positive.
Inside the table, every field except the RORI value can be edited. When editing the *Delta TTC* or *Annual Cost* field, the RORI of the corresponding CoA will be updated.
The table can be sorted by decreasing or increasing value for each column by clicking on the arrow next to the column name.

The *Hide* and *Show* buttons are used to reveal more details about each CoA. The table grows to show the defenses list of each CoA (Ref, Defense Name, Defense Info, Mitre Ref from the CoA Generator) and also a list of hosts impacted by the CoA.

Figure 3-4 displays the graphical interface of the RP with the full details of the CoAs.

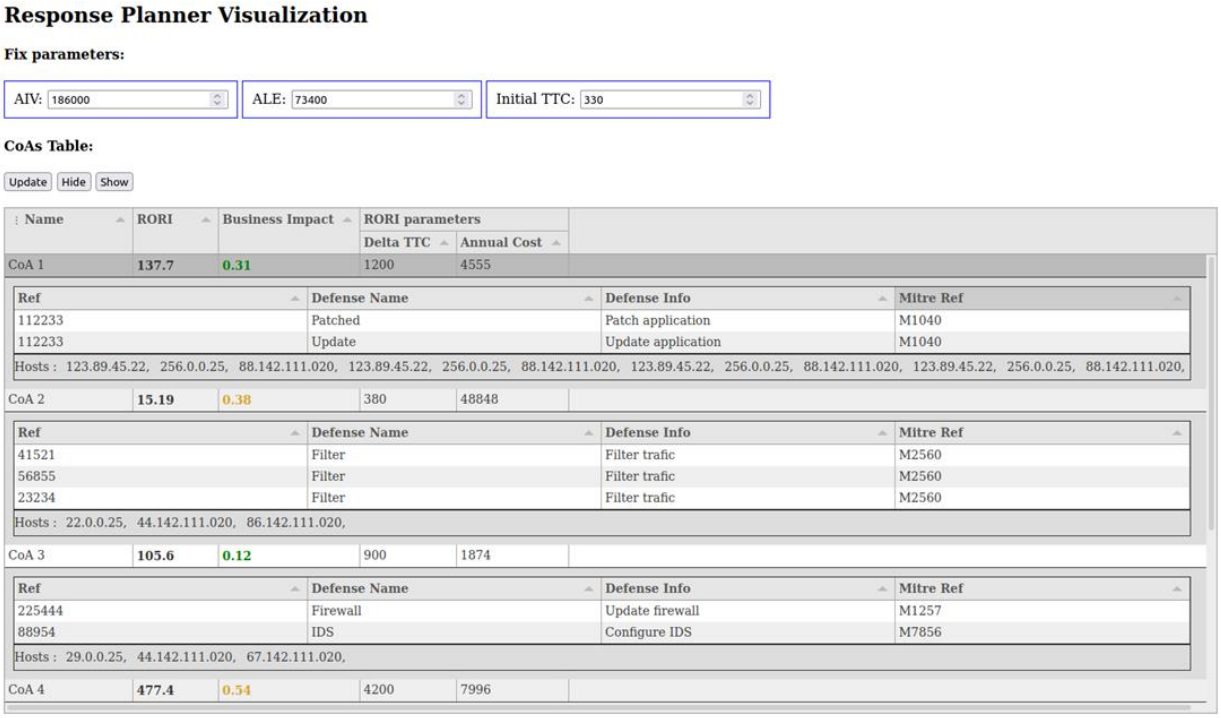

Figure 3-4 – Response Planner Graphical User Interface, full CoA display

The table inside the Web GUI is built using the JavaScript library *Tabulator*[7].

3.3 Next steps

In its current state, the Response Planner can compute the RORI metric and display the results on an interactive graphical interface. The next steps of the development will focus on the following points:

- Definition of the request to the BIA to compute the business impact of a CoA
- Definition of a way to evaluate the parameter *Annual Loss Expectancy* of the RORI metric. The estimation of this parameter should be based on the attack type deduced from the SIEM label, the criticality of the compromised assets computed by the BIA, and the initial TTC
- Generation of the containment CoAs according to the strategy described in Section 3.1.1.1
- Generation of a CACAO playbook associated to each CoA
- Update of the graphical interface to take into account the containment CoA, that will have no associated RORI value

[7] https://github.com/olifolkerd/tabulator

4 References

[CACAO] OASIS Collaborative Automated Course of Action Operations, Security Playbooks Version 1.0, available at https://docs.oasis-open.org/cacao/security-playbooks/v1.0/security-playbooks-v1.0.html

[D2.2] Deliverable D2.2, "System architecture & Interface specification; Initial version", SOCCRATES project

[D2.3] Deliverable D2.3, "System architecture & Interface specification; Final version", SOCCRATES project

[D5.1] Deliverable D5.1, "Definition of the business model structure, inputs and interfaces", SOCCRATES project

[OpenC2] Open Command and Control standard, available at https://openc2.org/

5 Abbreviations

This glossary serves as inventory of abbreviations used in the document.

Acronym	Description
AAD	AI-based Anomaly Detection
ADG	Attack Defence Graph
AI	Artificial Intelligence
AIT	AIT Austrian Institute of technology
AIV	Annual Infrastructure Value
ALE	Annual Loss Expectancy
API	Application Programming Interface
ARC	Annual Response Cost
BC	Business Company
BIA	Business Impact Analyser
BPMN	Business Process Model and Notation
CACAO	Collaborative Automated Course of Action Operations
CERT	Computer Emergency Response Team
CSIRT	Computer Security Incident Response Team
CoA	Course of Action
ICT	Information and Communication Technology
IMC	Infrastructure Modelling Component
IMT	Institut Mines Télécom – Télécom SudParis
IP	Internet Protocol
OIE	Orchestration & Integration Engine
OpenC2	Open Command and Control
RM	Risk Mitigation
RORI	Return on Response Investment
RP	Response Planner
SIEM	Security information and event management
SOAR	Security Orchestration, Automation and Response
SOC	Security Operation Centre
SOCCRATES	SOC & CSIRT Response to Attacks & Threats based on attack defence graph Evaluation Systems
TTC	Time To Compromise
UC	Use Case