SOCCRATES

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

D5.4

Evaluation and validation report

Deliverable type:	R
Contributing work packages:	WP5 Impact Quantification
Due date of deliverable:	31/10/2022
Submission date:	31/10/2022
Dissemination level:	PU

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

Abstract	This deliverable describes the evaluation and validation outcome of the Impact Analyser and Response Planner component on the use cases and experimentation platforms of the SOCCRATES project
Keywords:	Business Logic Modelling, Business Impact Analyser, Response Planner, Containment CoA Generator, Evaluation, Validation

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 Response 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.

This deliverable provides an evaluation and validation report on these two components. For the Business Impact Analyser, the report focuses on the relevancy of the results as well as the scalability of the computations. For the Response Planner, the report focuses on the coherence of the RORI metric as well as on its relevancy through use case studies.

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

Reviewer(s)	Frank Fransen (TNO)

Security Assessment	**See deliverables table for security assessment requirements**
Approval Date	27/10/2022
Remarks	-

TABLE OF CONTENTS

1 INTRODUCTION – RATIONALE OF THIS DOCUMENT 5
1.1 The SOCCRATES project 5
1.2 This deliverable 6
2 BUSINESS IMPACT ANALYSER 7
2.1 Business Logic Model generation 7
2.2 Assessment of the functionalities of the Business Impact Analyser 8
2.2.1 Evaluation of asset criticality computation 9
2.2.2 Evaluation of impact probability computation 10
2.2.3 Critical time computation 11
2.2.4 Summary of the assessments 12
2.3 Scalability of the computations 13
2.4 Conclusion 14
3 RESPONSE PLANNER 15
3.1 Evaluation of the RORI metric 15
3.1.1 RORI: financial impact quantification metric 15
3.1.2 RORI formula limits 17
3.1.3 Case studies: influence of parameter variations 18
3.2 Final updates of the Response Planner 21
3.2.1 CACAO Playbooks 21
3.2.2 Graphical user interface 21
3.3 Conclusion 22
4 REFERENCES 23
5 ABBREVIATIONS 24

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.

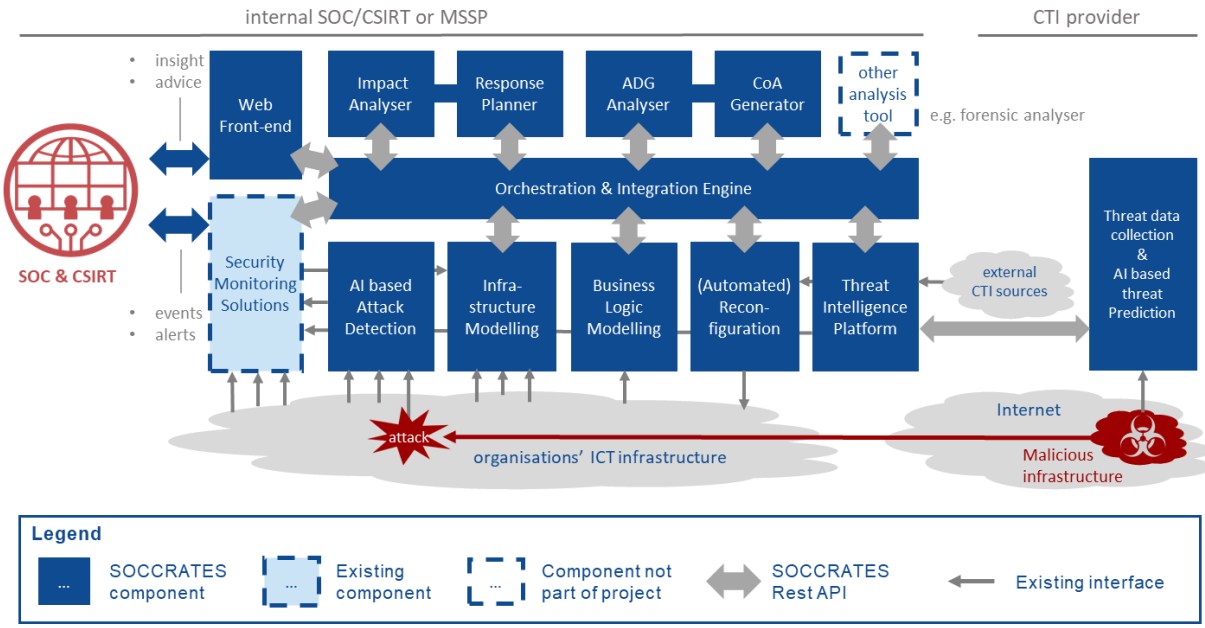

Figure 1.1 – The SOCCRATES platform

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 provide evaluation and validation data for the two components developed by IMT for the SOCCRATES project, namely the Business Impact Analyser and the Response Planner. The theoretical background of the Business logic modelling integrated in the Business Impact Analyser was already presented in Deliverable D5.1 [D5.1], while the description of the final prototypes of these two components was provided in Deliverable D5.3 [D5.3].

The approach for evaluation and validation depends on the component. For the Business Impact Analyser, both the relevancy of the results and the scalability of the computations are investigated, while for the Response Planner, the relevancy of the results and the influence of parameter variation are considered. Indeed, the Business Impact Analyser relies on a graph model in which computations depend on the number of nodes and edges, entailing scalability issues that require attention. The Response Planner, in contrast, does not rely on a complex model, and therefore does not present any scalability issue. However, since the Response Planner evaluates and ranks CoAs based on the RORI metric, a study on the parameter variations provides useful insight regarding the relevancy and the validation of the metric.

1.3 Structure of this deliverable

Section 2 of this deliverable describes the evaluation and validation of the Business Impact Analyser, first regarding the relevancy of the outputs, and then regarding the scalability of the component. Section 3 describes the evaluation and validation of the Response Planner, first regarding the relevancy of the outputs, and then regarding the influence of parameter variation.

2 Business Impact Analyser

This section presents a simple model for the Business Impact Analyser (BIA), which will make it possible to evaluate and validate its different functionalities:

- The calculation of the impact on the business entities : impact probability and critical time
- The computation of the asset criticality

First, the BIA needs to build a model. To generate this model, the BIA relies on the Business Logic Modelling, which builds a mission dependency model and a resource dependency model. These two models are then fused into a matrix that represents the business logic model, as detailed in D5.1 [D5.1] and D5.3 [D5.3]. The purpose of designing a simple model is to give intuitive explanations regarding the computations, making it easy to check that the output values are consistent with the model.

2.1 Business Logic Model generation

The evaluation scenario is based on a graph comprised of the following nodes:

- 6 Assets: A, B, C, D, E and F.
- 3 Business Functions: BF1, BF2 and BF3
- 2 Business Processes: BP1 and BP2
- 1 Business Company: BC

The business logic model is based on two adjacency matrixes:

- Adjacency matrix with conditional impact probabilities
- Adjacency matrix with downtime tolerances

Table 2-1 presents the adjacency matrix with conditional impact probabilities which is built from these two matrixes: the Business Entity probability matrix and the Inter-assets probability matrix. The generation of this matrix is described in more detail in D5.3 [D5.3].

Table 2-1 - Adjacency matrix with conditional impact probabilities

node	A	B	C	D	E	F	BF1	BF2	BF3	BP1	BP2	BC
A	0.0	0.46	0.52	0.0	0.21	0.27	0.51	0.0	0.0	0.0	0.0	0.0
B	0.04	0.0	0.0	0.54	0.0	0.26	0.0	0.78	0.0	0.0	0.0	0.0
C	0.22	0.0	0.0	0.0	0.62	0.0	0.24	0.0	0.0	0.0	0.0	0.0
D	0.0	0.39	0.0	0.0	0.0	0.29	0.0	0.8	0.85	0.0	0.0	0.0
E	0.43	0.0	0.48	0.0	0.0	0.18	0.0	0.0	0.0	0.0	0.0	0.0
F	0.31	0.15	0.0	0.46	0.17	0.0	0.0	0.0	0.0	0.0	0.0	0.0
BF1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.49	0.0	0.0
BF2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.62	0.64	0.0
BF3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.51	0.0
BP1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.72
BP2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.81
BC	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

The adjacency matrix with the downtime tolerances is represented by **Table 2-2**. It is built from two matrixes: the Business Entity downtime tolerance matrix and the inter-asset downtime tolerance matrix.

Table 2-2 - Adjacency matrix with downtime tolerance

node	A	B	C	D	E	F	BF1	BF2	BF3	BP1	BP2	BC
A	0.0	0.1	0.1	0.0	0.1	0.1	1.2	0.0	0.0	0.0	0.0	0.0
B	0.1	0.0	0.0	0.1	0.0	0.1	0.0	2.3	0.0	0.0	0.0	0.0
C	0.1	0.0	0.0	0.0	0.1	0.0	2.5	0.0	0.0	0.0	0.0	0.0
D	0.0	0.1	0.0	0.0	0.0	0.1	0.0	1.5	1.85	0.0	0.0	0.0
E	0.1	0.0	0.1	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
F	0.1	0.1	0.0	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
BF1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	3.2	0.0	0.0
BF2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.9	2.1	0.0
BF3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.65	0.0
BP1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	2.32
BP2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.74
BC	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

2.2 Assessment of the functionalities of the Business Impact Analyser

Figure 2-1 presents the graphical interface of the Business Impact Analyser with the probability graph, is built from the adjacency matrix with conditional impact probabilities provided in Table 2-1.

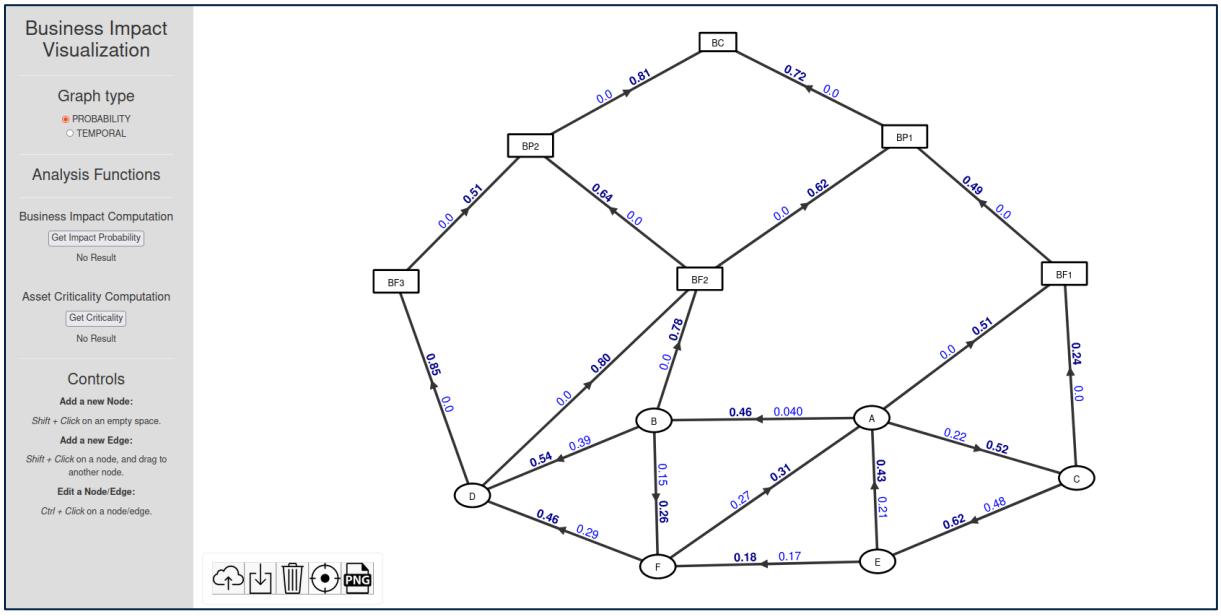

Figure 2-1 - Probability graph

Similarly, **Figure 2-2** displays the temporal graph which is built from the adjacency matrix with the downtime tolerance provided in **Table 2-2**.

Classification level: Public

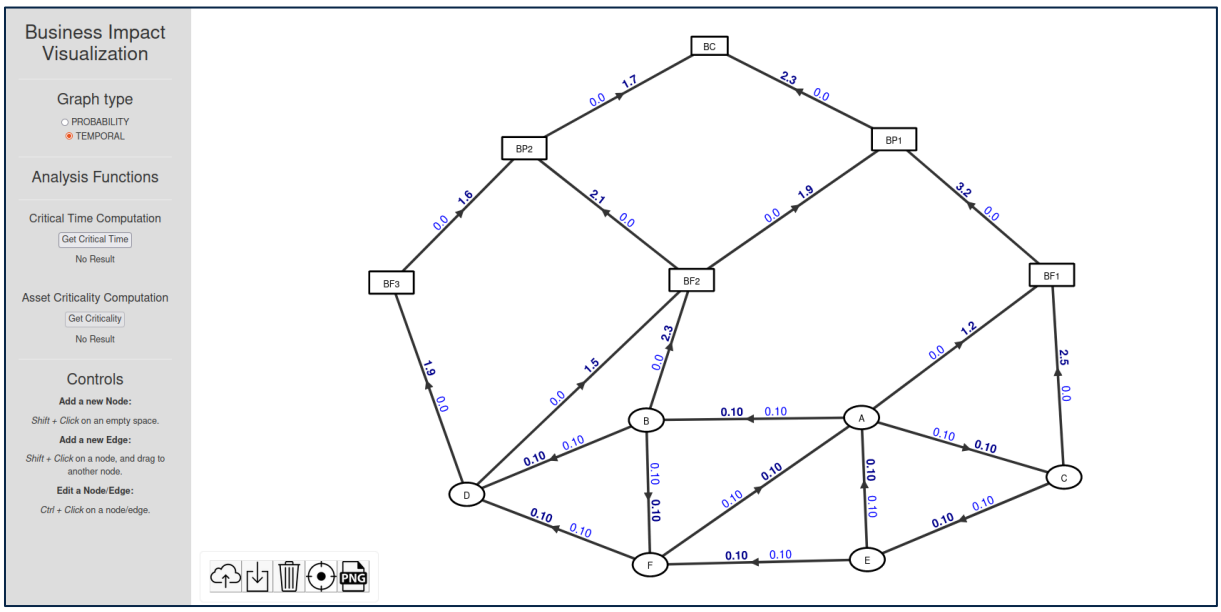

Figure 2-2 - Temporal graph

2.2.1 Evaluation of asset criticality computation

In order to identify the assets that contribute the most to the propagation of impacts on the business entities, the Business Impact Analyser calculates the criticality of the assets, and based on the computed values, it classifies the assets into three levels of criticality: high, medium and low.

In order to identify the assets that contribute the most to the propagation of impacts on the business entities, the Business Impact Analyser calculates the criticality of the assets, and based on the computed values, it classifies the assets into three levels of criticality: high, medium and low.

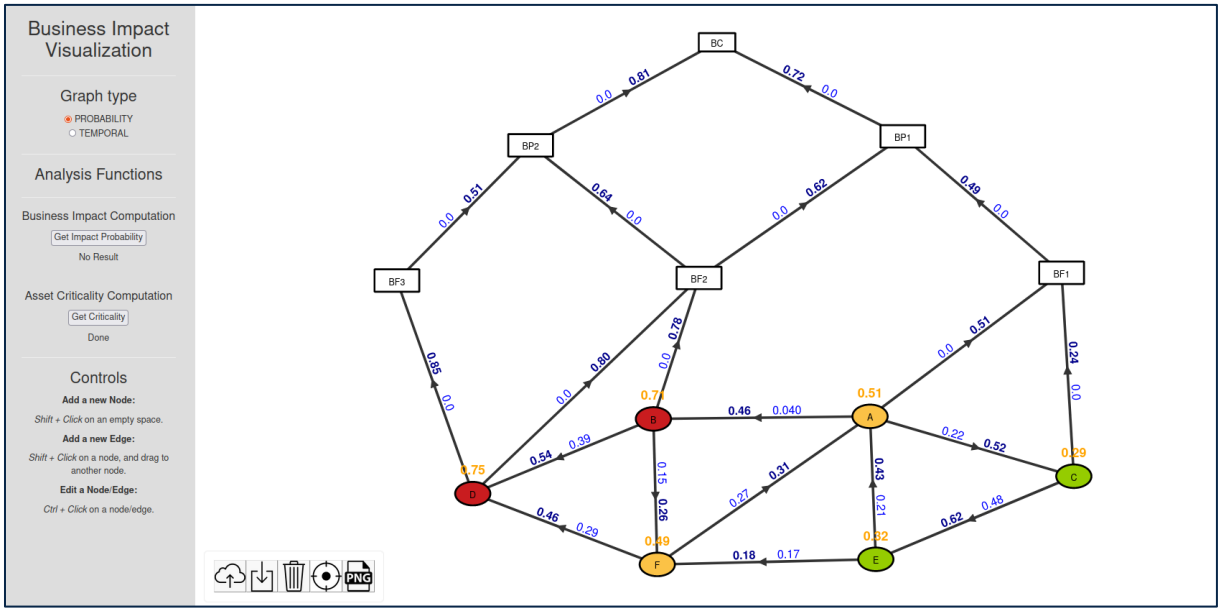

Figure 2-3 - Asset criticality computation

Figure 2-3 displays the graphical interface with the probability graph after the computation of the asset criticality.

In order to quickly identify the most critical and vital assets in this business logic model, the Business Impact Analyser calculated the criticality of assets in 0.052 seconds, and displayed the assets in three colours according to the criticality values, as shown in **Figure 2-3**:

- In green : Assets C and E are considered as low critical assets
- In orange : Assets A and F are considered as medium critical assets
- In red : Assets B and D are considered as highly critical assets

These results can be interpreted in the following way: asset B and D directly supports business functions BF2 and BF3, with high impact probabilities (all superior to 0.7). The impact on these business functions can then propagate to business processes BP1 and BP2, and then to the business company BC with still fairly high probabilities (all superior to 0.5). Therefore, an attack targeting these two assets would have the highest impact on most business entities and on the company as a whole, which explains why they are evaluated as critical.

In comparison, asset C, while supporting business function BF1, has a low impact probability on that entity (0.24). Moreover, in order to propagate the impact to business functions BF2 and BF3, the propagation must reach one or several intermediate assets (A, E and F), and these inter-assets conditional impact probabilities are all fairly low. Therefore, an attack on asset C is not likely to cause any serious disruption on the business entities and the company, which explains why this asset is not evaluated with a low criticality value.

The medium criticality score of assets A and F is due to their closeness to critical assets B and D. For both A and F, the impact of an attack could propagate to respectively B and D with a probability close to 0.5, making these assets not as critical as B and D, but still of medium importance.

Finally, since asset E is not close to any critical asset and has fairly low impact propagation probabilities on its neighbours, its criticality value is rated as low.

Based on this simple model, it is therefore possible to assess the relevancy of the computed criticality values, and understand the reasons for which the criticality of an asset will be rated as low, medium, or high.

2.2.2 Evaluation of impact probability computation

Figure 2-4 displays the computation of the impact probability of a shock event on the business entities. The shock event is given as input: (Name of the shock event, target asset, impact probability). In this graph, it is presented by: (Shock Event 1, Asset D, 0.80).

The Business Impact Analyser computes the probability of impact of this shock event on the business entities in 0.035 seconds, and displays the business entities nodes in three colours according to the calculated business impact:

- In green : the business impact on BF1 is low
- In orange : the business impact on BP1, BP2 and BC is moderate
- In red : the business impact on BF2 and BF3 is high

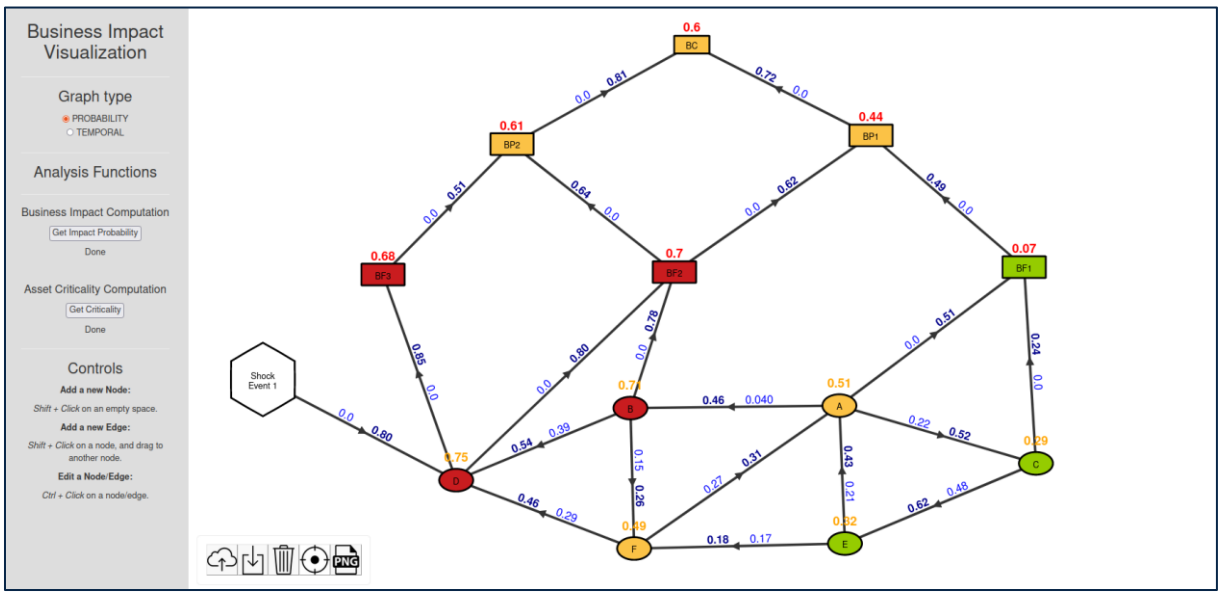

Figure 2-4 - Impact probability computation

The results of this computation can be interpreted as follows: overall, the farther a business entity is located from the target node, the lower the impact is. Obviously, this statement must be nuanced by taking into account the impact propagation probabilities on each edge, but this is still a reasonable approximation of the situation.

Asset D directly propagates the impact on business functions BF2 and BF3, both with high probabilities, which results in a high impact on these two business functions. Business function BF1 however can only be impacted of the initial event propagates though several intermediate nodes, which happens with low probabilities, and therefore results in a low impact on BF1.

The impact on business processes BP1 and BP2 is medium due to the impact probabilities between the business functions and business processes being only around 0.6, yet the impact on BP2 is significantly higher than on BP1. This is explained by the fact that BP2 receives the impact propagation from BF2 and BF3, which are both highly impacted by the shock event, while the impact contribution on BP1 comes from BF1 and BF2, and BF1 only has a very slight impact from the event.

Finally, the impact on the whole company is rated as medium with a value of 0.6. The reason why the impact is not rated as high, despite the shock event targeting a critical asset, is because the impact probability of this event was set at 0.8 instead of 1.

This analysis of the output values therefore shows that the impact probability computation on business entities following a shock event targeting an asset in the infrastructure is relevant, and is able to provide a SOC analyst with a global overview of the disturbance on the various business entities following an event such as a successful attack on an asset.

2.2.3 Critical time computation

The computation of the critical time was applied on the adjacency matrix with downtime tolerance using the Dijkstra algorithm.

Figure 2-5 displays the critical time value and the shortest path from the shock event to the business company.

The Business Impact Analyser displays the shortest path as a dotted line and calculated the critical time value, which is equal to the sum of the downtime tolerances between the nodes in the shortest path in 0.035 seconds.

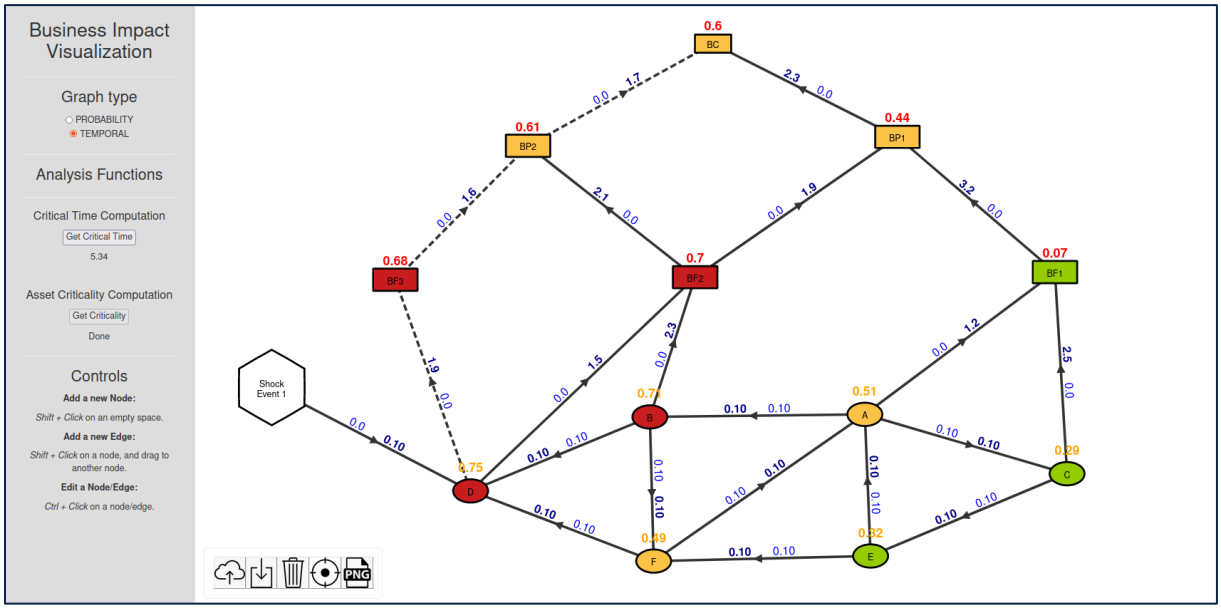

Figure 2-5 - Critical time computation

Since the computed critical time is only the result of a shortest path algorithm, no further analysis is needed here to assess the relevancy of the result. However, it should be emphasized that an extension of this model, implemented aside from the SOCCRATES project, takes into account several sets of downtime tolerance values to account for various situations that depend on the time of day, such as the difference in activity between day time and night time.

2.2.4 Summary of the assessments

Table 2-3 summarizes the outputs of the Business Impact Analyser, classified by functionalities.

Table 2-3 – Summary of the BIA outputs

Functionalities	Results	Computation time (in sec)
Asset criticality computation	3 levels of criticality : • Highly critical assets : Asset B and Asset D. • Medium critical assets : Asset A and Asset F. • Low critical assets : Asset C et Asset E.	0.052
Impact probability computation	3 levels of impact of this shock event : (Shock Event 1, Asset D, 0.80) on the business entities :	0.035

Classification level: Public

	• High impact : Business Function 2 and Business Function 3. • Moderate impact : Business Process 1, Business Process 2 and Business Company. • Low impact : Business Function 1	
Critical time computation	The BIA returned : • Critical time : 5.34 hours • Shortest path : Shock Event 1 → Asset D → Business Function 3 → Business Process 2 → Business Company	0.035

2.3 Scalability of the computations

This section presents the scalability of the computations performed by the Business Impact Analyser component. In order to assess the scalability of the computations, several tests were performed with different graphs, with different number of assets and edges.

Table 2-4 presents the different tests performed and their results. Each line of the table presents a test, with the size of the graph and the time needed to compute the criticality of the assets. Three colours are used to distinguish the edge density in the graph:

- black: all assets communicate with each other, which means that all assets send and receive data (complete graph)
- red: half of the assets communicate with each other, and the rest of the assets only receive data.
- green: a quarter of the assets communicate with each other, and the rest of the assets only receive data.

Number of assets	Number of edges	Computation time (in sec)	Computation time / asset (in sec)
50	2 450	1.96	0.039
50	1 224	1.09	0.021
50	613	0.69	0.013
250	62 250	24.0	0.095
250	31 124	12.2	0.048
250	15 562	6.45	0.025
1 000	999 000	197.6	0.197
1 000	499 500	101.9	0.101
1 000	249 750	52.4	0.052

Table 2-4 - Tests of scalability of computations

The chart in **Figure 2-6** shows the dependency between the computation time and the number of edges. It is important to note that the computation performed in this test is the criticality value of all the assets, which is the longest type of request that can be sent to the BIA since it requires to compute the value for all the nodes in the graph. That is why **Table 2-4** also provides a column giving the computation time per asset, which corresponds to the impact probability request following a shock event targeting a specific asset.

These results show that the BIA is able to make computations in reasonable time in fairly large graphs (such as complete graphs with 250 nodes). In graphs comprised of a 1000 assets, the computation time starts to reach values that induce a substantial delay, but the computation still remains feasible.

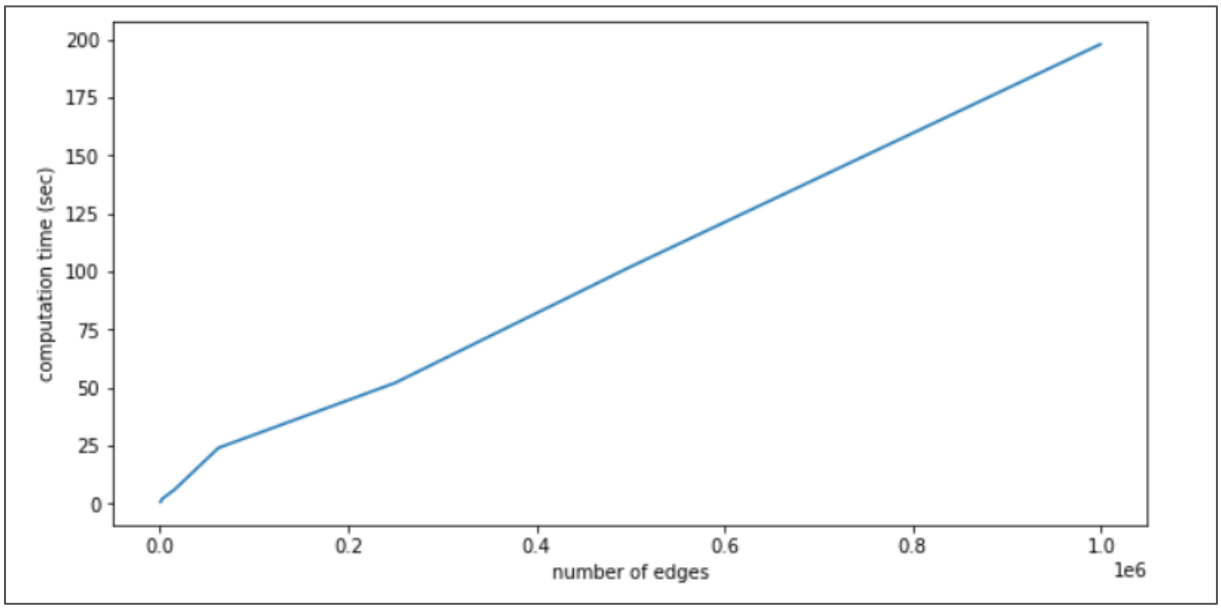

Figure 2-6 - Dependency between the computation time and the size of the graph

2.4 Conclusion

The current version of the Business Impact Analyser produces relevant results in very reasonable times for fairly large graphs. Even though companies might have more than 1000 assets in their infrastructure, in practice, and based on the model built at Mnemonic, a large number of assets do not regularly exchange data, which results in a lot of isolated nodes which do not add any complexi-ty to the computation. Therefore, the prototype of the Business Impact Analyser can be deployed even in fairly large companies to provide real time business impact assessment.

3 Response Planner

3.1 Evaluation of the RORI metric

One functionality of the Response Planner (RP) is to generate an evaluation of the CoAs, in order to rank, compare and select the most suitable CoAs to deploy. Depending on the type of CoAs, the RP performs an evaluation with an appropriate metric. For classical CoAs, since they are deployed to obtain a long term benefice, a return on response investment (RORI) metric is more appropriate. On the other hand, containment CoAs are expected to be deployed temporarily as short-term countermeasures. Therefore, an operational impact quantification metric is more relevant to evaluate the immediate impact of the Containment CoA on company business processes.

The analysis of the business impact metric is already provided in Section 2. Therefore, this section focuses on detailing the evaluation of the RORI metric relevancy. First, thorough analysis of the RORI formula is provided, with the influence of parameter variation and the study of the formula at its limits, and then, the RORI value results are provided on a realistic use case. Contrary to the BIA, there is no need to study the scalability for the RP since there is no model building, and therefore no specific complexity in the computations.

The tricky part of computing the RORI value in the context of the SOCCRATES project is that the RORI evaluates a financial impact on the company and as such, requires financial data in its parameters, which were not available to our partners. Therefore, the estimation of the RORI parameters has been adapted to match the available data from the other components, mainly the CoA Generator and the ADG.

3.1.1 RORI: financial impact quantification metric

As presented in D5.3 [D5.3], the RORI metric is based on four parameters: Annual Loss Expectancy (ALE), Risk Mitigation (RM), Annual Infrastructure Value (AIV), Annual Response Cost (ARC), and is computed as follows:

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

The estimation of the Annual Loss Expectancy and Risk Mitigation parameter has been updated since D5.3 [D5.3]. Indeed, the high-value assets risk is now fetched directly from the Attack Defense Graph (ADG) simulations, instead of estimating it from the Time to Compromise (TTC) of assets. The detailed formulae to calculate the value of the ALE and the RM parameters are provided in Section 3.1.1.1 and Section 3.1.1.2 respectively.

The RORI of a CoA has to be positive for the CoA to be profitable to the company.

3.1.1.1 Annual Loss Expectancy (ALE)

The Annual Loss Expectancy parameter is the expected loss from the attack if no CoA is deployed. It is dependent on two parameters, Single Loss Expectancy (SLE) and Annual Rate of Occurrence (ARO).

$$ALE = SLE \times ARO$$

The Single Loss Expectancy is the monetary loss we expect each time the attack occurs. It is the multiplication of the Exposure Factor (EF) by the monetary value estimation of the impact on high-value assets.

$$SLE = [\text{Impacted_assets}_{\text{monetary-value}} \times EF]$$

The Exposure Factor (EF) is the percentage of the value of a given asset that gets lost as a result of an incident. A loss can exceed the value of a given asset; in such cases, the EF would be greater than 1. EF could be determined for each asset impacted, but it have been simplified to a single value of EF for a threat. Therefore, an exposure factor value has been set for each event type (obtained from the AI-based Attack Detection (AAD) component). The choice that was made for each event type is as follows:

Event type	Initial Access	Lateral Movement	Data Exfiltration	Others
Exposure Factor	0.1	0.4	0.8	0.3

$$\text{Impacted_assets}_{\text{monetary-value}} = \frac{risk_{initial}}{risk_{max}} \times AIV$$

In the SOCCRATES platform, the monetary value of the impact on high-value assets is estimated by computing the ratio of the initial risk over maximum risk, and multiplying it by the Annual Infrastructure Value (AIV). The initial risk and maximum risk are parameters computed by the ADG component. The initial risk represents the amount of consequence that the company is exposed to, by the threat under investigation, without any CoA deployed. The maximum risk represents the maximum risk value if all the high-value assets were entirely compromised. These risk estimations are expressed in a unit that is only relevant for the ADG, however the ratio of the initial risk over maximum risk allows the RP to estimate the percentage of the company that is at risk.

The Annual Rate of Occurrence (ARO) is the number of times the vulnerability is likely to be exploited in a year. It is the product of the Annual Attack Rate and the Exploit Probability.

$$ARO = AAR \times EP$$

The Annual Attack Rate (AAR) is the estimation of the number of attacks that the company faces in one year. This estimation should be given by the company.

The Exploit Probability (EP) is the probability that an attacker will exploit the threat vulnerability. The SOCCRATES platform does not have enough data yet to predict such a probability, therefore it has been set to 1.

To sum it up, the Annual Loss Expectancy is obtained using the following formula:

$$ALE = SLE \times ARO = \frac{risk_{initial}}{risk_{max}} \times AIV \times EF \times AAR$$

3.1.1.2 Risk Mitigation (RM)

The Risk Mitigation is now estimated as follows:

$$\text{Risk Mitigation} = [1 - \frac{risk_{coa}}{risk_{initial}}]$$

If $risk_{coa} \geq risk_{initial}$, the risk is higher after applying the CoA. In such situations, we consider that the CoA is harmful, and we set RM equal to 0.

If $risk_{initial} = 0$, there is no initial risk, then there is nothing to mitigate. We set RM equal to 0.

Based on this formula, it is possible to check that *RM* approaches 1 when *risk_coa* tends to 0, which reflects the fact that the risk is entirely mitigated. Conversely, when *risk_coa* is close to *risk_initial*, *RM* approaches 0, since the CoA does not improve the situation.

3.1.2 RORI formula limits

3.1.2.1 Maximum RORI value

In this case study, it is supposed that the Exposure Factor, the Annual Attack Rate and the initial risk are fixed. In order to evaluate the maximum RORI value that a CoA could reach, the Annual Response Cost has to be negligible and the Risk Mitigation has to be equal to 1. It is the upper limit, where the initial risk is totally mitigated by the CoA and there is a negligible response cost.

$$RORI_{max} = \frac{ALE}{AIV} \times 100$$

With the parameter estimation used in the SOCCRATES project, it can be written as:

$$RORI_{max} = \frac{risk_{initial}}{risk_{max}} \times EF \times AAR \times 100$$

For a maximal initial risk, (risk_initial = risk_max), the maximal RORI value is equal to the Exposure Factor, times the Annual Occurrence of the threat.

3.1.2.2 Upper limit of the RORI for a CoA:

For a CoA with a fixed Risk Mitigation, the limit to its RORI will be:

$$\text{RORI_CoA}_{max} = RM_{coa} \times \frac{ALE}{AIV} \times 100 = RM_{coa} \times RORI_{max}$$

3.1.2.3 Lower limit of the RORI

To approach the RORI lower limit value, the Risk Mitigation has to tends to 0. It is the situation where the initial Risk is not mitigated at all by the CoA.

$$\lim_{\text{risk_coa} \to \text{risk_initial}} RM = 0$$

$$\lim_{RM \to 0} RORI_{max} = 0$$

$$\lim_{RM \to 0} RORI = \frac{-ARC}{ARC + AIV} \times 100$$

The lower RORI limit is therefore negative, and mainly depends on the Annual Response Cost value.

3.1.2.4 RORI sign

The RORI is positive if the Annual Loss Expectancy times the Risk Mitigation is bigger than the Annual Response Cost.

$$RORI > 0 \; if \; ALE \times RM > ARC$$

3.1.3 Case studies: influence of parameter variations

3.1.3.1 Maximum budget for a CoA with a defined Annual Loss Expectancy

Figure 3-1 displays the evolution of the RORI of a CoA with a complete risk mitigation (RM=1), according to its Annual Response Cost, with different ALE values. The AIV of the company is estimated to be 2 000 000€.

It can be observed that the more expensive a CoA, the lower its RORI. While the CoA RORI is positive, it remains cost-effective for the company to deploy the CoA. Therefore, the maximal budget for the response can be estimated at the point where the RORI is no longer positive.

For example, if a threat has an ALE of EUR 512 000, the considered CoAs must have a cost below EUR 512 000 in order to have a positive RORI and therefore be profitable for the company. This budget limit will be lower if the CoA has a lower risk mitigation. It is also worth noting that the higher the ALE, the higher the RORI values, which can be explained by the fact that a higher ALE value indicates that the attack would have a higher financial impact on the company, therefore increasing the profitability for the company to deploy a given CoA.

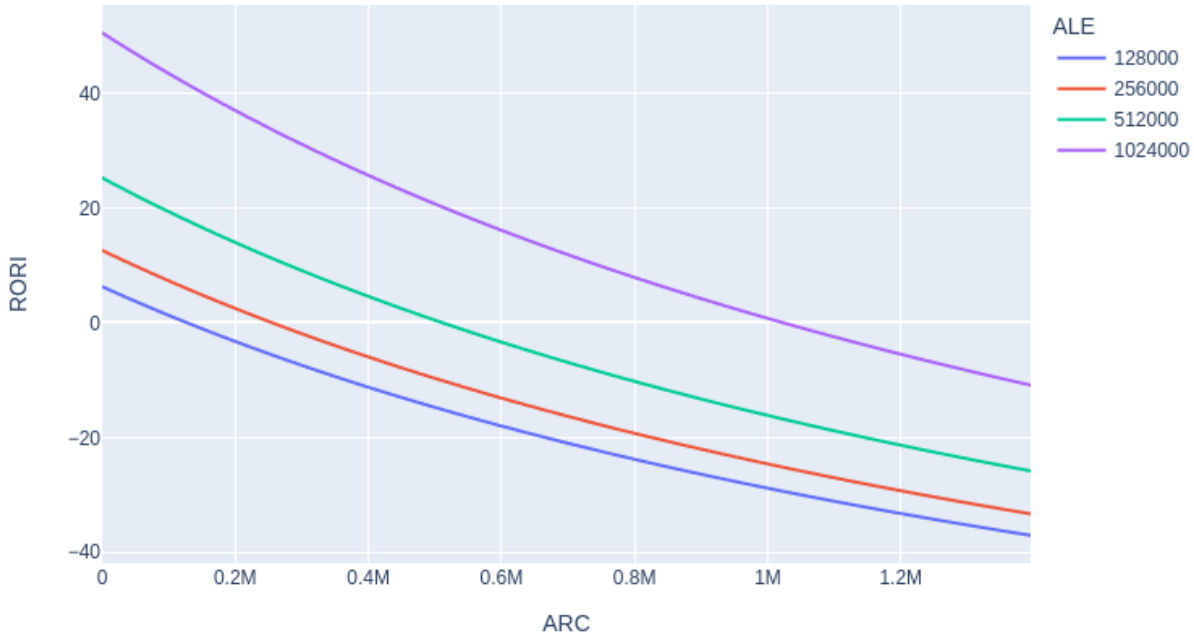

Figure 3-1 – RORI values variation with the Annual Response Cost, for several ALE values

3.1.3.2 Maximum budget for a CoA with a defined risk mitigation

This second use case considers the same company AIV as in the first use case, 2 000 000€. The threat is predetermined as a Lateral Movement inside the company network, having eighty percent of its business functions impacted by it, and may happen two times a year. From these hypotheses, the Annual Loss Expectancy is estimated at 1,28 million euros. From these values, the maximum RORI is 64%.

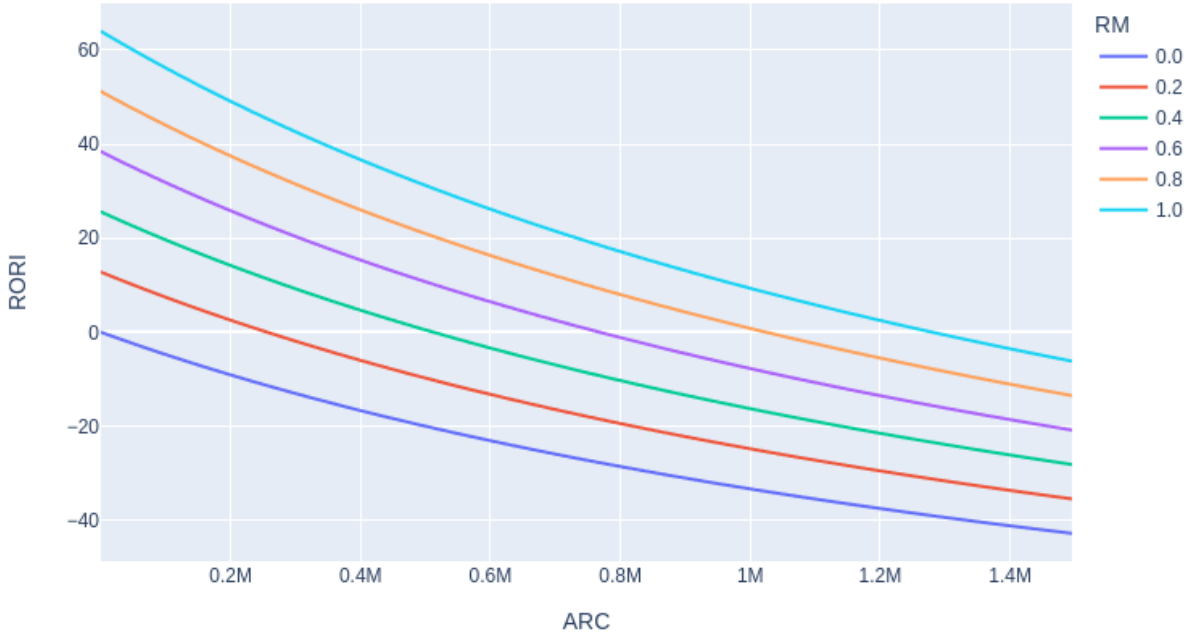

Figure 3-2 – RORI values for different CoA Risk Mitigation, according to the Annual Response Cost

Figure 3-2 displays the evolution of the RORI depending on the CoA cost, for various RM values. For example, the company could consider a CoA to help in the detection of prevention of Lateral Movement, such as mapping lateral movement paths and using an additional reporting tool. The company could aim for this countermeasure to help to reduce the risk of the initial threat by sixty percent. Then using **Figure 3-2**, the company is able to figure out that this CoA will be profitable if its cost is below 700 000€.

3.1.3.3 Comparison between two CoAs

This use case considers the same type of malicious activity as in the previous use case, Lateral Movement.

Two CoAs have been proposed to mitigate the threat. The first one is rather accessible (EUR 15 000) and have an average risk mitigation of 30%.
The second CoA is more expensive (EUR 300 000) but has a higher risk mitigation of 80%.

The company is confronted to a dilemma, either chose the cheapest CoA but the company could be more exposed in the future, or chose to invest in the second CoA, more expensive but with a better coverage, hoping to get a better return on response investment by avoiding more threats over the long term.

The answer to this dilemma depends on the size of the company, and its Infrastructure Value. The more expensive the infrastructure is, the more cost-effective it is to invest in its protection.

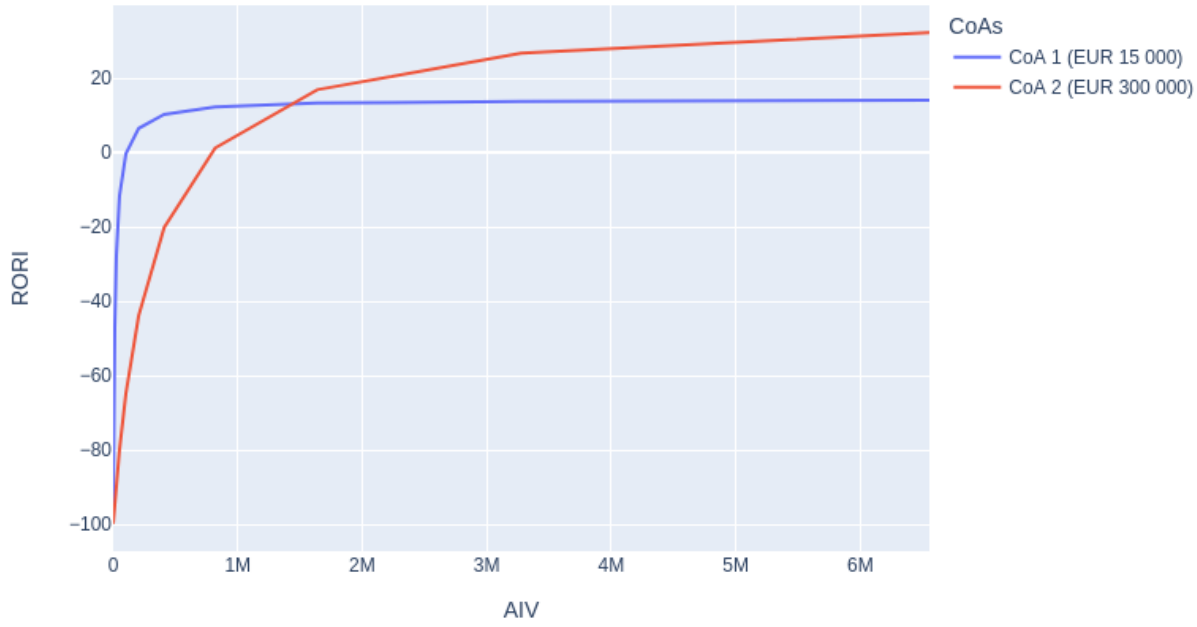

Figure 3-3 – Comparison of RORI values for two CoAs, according to the Annual Infrastructure Value

Thanks to

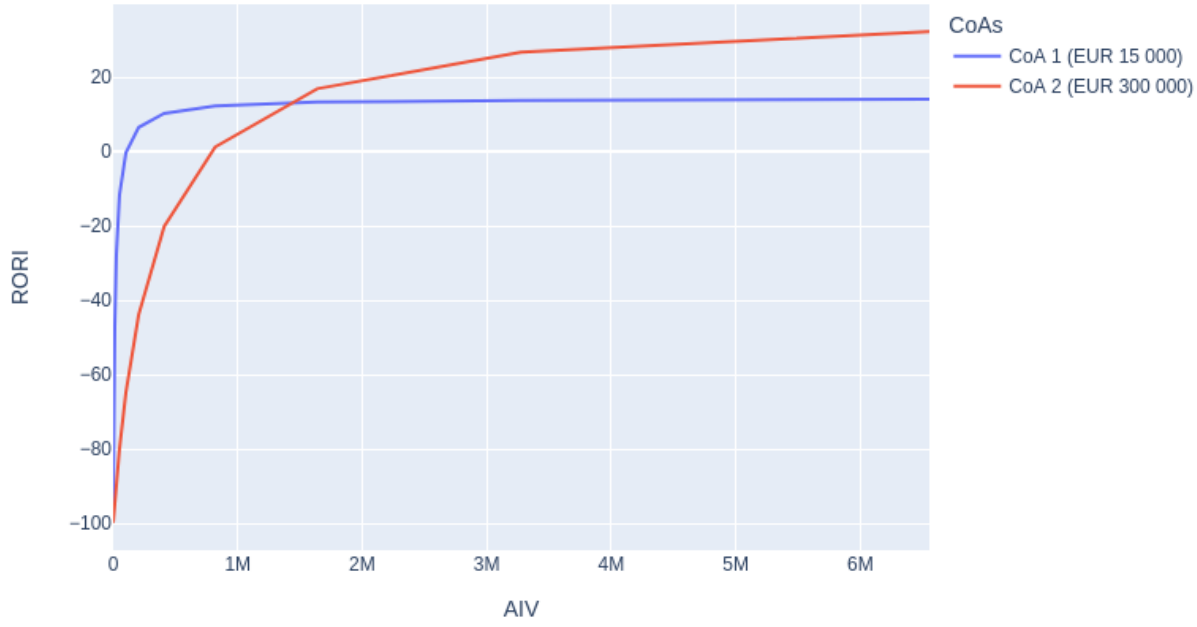

Figure 3-3, several company size cases can be studied. It can be observed that for a small company, with a low AIV, the first CoA is the best choice. Indeed, the first CoA, due to its low cost, rapidly reaches a positive RORI value, meaning its deployment will be profitable to the company. Whereas, the second CoA has a negative RORI value, meaning that the deployment cost is too high to be balanced by benefits of the threat mitigation.

For a company with an AIV around 1,5 million euros, the two CoAs are evaluated with similar RORI value, meaning that the two CoAs are equally profitable to invest in.

For a big company with an AIV superior to four million euros, the second CoA becomes largely more profitable. The upper RORI limit of the first CoA is limited 14%, due to its low risk mitigation. On the other hand, the investment to the second CoA keeps being more profitable as the AIV increase. The second CoA therefore reaches a higher RORI value of 38%.

3.2 Final updates of the Response Planner

This section briefly presents the latest updates made in the Response Planner. The features presented in this section do not require any specific evaluation or validation.

3.2.1 CACAO Playbooks

The generation of OpenC2 [OpenC2] commands inside CACAO playbooks [CACAO] for Containment CoAs has been implemented in the Response Planner. It is now possible to automatically deploy Containment CoAs, thanks to the appropriate OpenC2 commands.

3.2.2 Graphical user interface

The graphical user interface of the Response Planner has been updated, in order to be easily readable. Its final version is presented in **Figure 3-4**.

Figure 3-4 – Update of the RP frontend display

Classification level: Public

3.3 Conclusion

The current version of the Response Planner fulfills its two objectives: generation of Containment CoAs OpenC2 commands for automated deployment, and CoA ranking. The RORI metric used for the ranking of CoAs is adapted to the inputs available from the other components of the SOCCRATES project, and produces results that can be easily interpreted and explained. However, a company that deploys the Response Planner would be free to do the parameter estimation differently from the proposed approach, for example by integrating financial data that were not available during the project.

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

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

[D5.3] Deliverable D5.3, "Business Logic Modelling and Impact Analyser & Response Planner – Final prototype", 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 Attack Detection
ACT	semi-Automated Cyber Threat intelligence
ADG	Attack Defence Graph
AEF	Argus Event Format
AI	Artificial Intelligence
AIT	AIT Austrian Institute of technology
AIV	Annual Infrastructure Value
ALE	Annual Loss Expectancy
API	Application Programming Interface
APT	Advance Persistent Threat
ARC	Annual Response Cost
ATOS	ATOS Spain
AV	AntiVirus
BC	Business Company
BF	Business Function
BIA	Business Impact Analyser
BP	Business Process
BPMN	Business Process Model and Notation
CACAO	Collaborative Automated Course of Action Operations
CC	Command and Control
CI/CD	Continuous Integration and Continuous Deployment[1]
CERT	Computer Emergency Response Team
CMDB	Configuration Management Database
CSIRT	Computer Security Incident Response Team
CoA	Course of Action
CTI	Cyber Threat Intelligence
DC	DataCentre
DGA	Domain Generated Algorithm
DNS	Domain Name System
EDR	Endpoint Detection and Response
ELK	Elasticsearch/Logstash/Kibana
FRS	Foreseeti
FSC	F-secure
ICT	Information and Communication Technology
IDS	Intrusion Detection System
IMC	Infrastructure Modelling Component
IMT	Institut Mines Télécom – Télécom SudParis
INTF	Interface
IoC	Indicators of Compromise

[1] It is common that the CD refers to Continuous Delivery. In this report, we prefer Continuous Deployment.

IP	Internet Protocol
IPS	Intrusion Prevention System
IRM	Incident Response and Management
ITIL	Information Technology Infrastructure Library
KTH	Kungliga Tekniska högskolan - Royal Institute of Technology
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
M_n	Infrastructure Model (at time *n*)
MNM	mnemonic
MSSP	Managed Security Service Provider
MTTD	Mean Time To Detection
NOC	Network Operations Centre
OpenC2	Open Command and Control
OT	Operational Technology
OS	Operating System
RM	Risk Mitigation
RORI	Return on Response Investment
RP	Response Planner
SDN	Software Defined Network
SHS	Shadowserver
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
SSL	Secure Sockets Layer
TAP	Test Access Point
TIP	Threat Intelligence platform
TLS	Transport Layer Security
TNO	Nederlandse Organisatie voor toegepast natuurwetenschappelijk onderzoek
TTC	Time To Compromise
UC	Use Case
VLAN	Virtual LAN
VM	Virtual Machine
VTF	Vattenfall