SENTINEL

Bridging the security, privacy, and data protection gap for smaller enterprises in Europe

D4.1-The SENTINEL services: MVP

This work is part of the SENTINEL project. SENTINEL has received funding from the European Union's Horizon 2020 Work Programme for research and innovation 2018-2020 under grant agreement n°101021659.

Project Information

Grant Agreement Number	**101021659**
Project Acronym	SENTINEL
Project Full Title	Bridging the security, privacy, and data protection gap for smaller enterprises in Europe
Starting Date	1st June 2021
Duration	36 months
Call Identifier	H2020-SU-DS-2020
Topic	H2020-SU-DS-2018-2019-2020 Digital Security and privacy for citizens and Small and Medium Enterprises and Micro Enterprises
Project Website	https://www.sentinel-project.eu/
Project Coordinator	Dr. George Bravos
Organisation	Information Technology for Market Leadership (ITML)
Email	gebravos@itml.gr

Document Information

Work Package	Work Package 4
Deliverable Title	D4.1 – The SENTINEL services: MVP
Version	1.2
Date of Submission	30/05/2022
Main Author(s)/ Editor(s)	Yannis Skourtis (IDIR), Evangelia Kavakli (IDIR), Peri Loucopoulos (IDIR)
Contributor(s)	Christos Dimou (ITML), Thomas Oudin (ACS), Michalis Smyrlis (STS)
Reviewer(s)	Philippe Valoggia (LIST), Nikos Nikolaou (ITML)

Document Classification							
Draft		**Final**	X	**Confidential**		**Public**	X

History			
Version	**Issue Date**	**Status**	**Distribution**
0.1	15/04/2022	Draft	Confidential
0.5	16/5/2022	Draft	Confidential
0.7	18/5/2022	Draft	Confidential
0.8	19/5/2022	Draft	Confidential
1.0	19/5/2002	Draft	Confidential
1.1	27/5/2022	Draft	Confidential
1.2	30/05/2022	Final	Public

Table of Contents

List of Figures

Abbreviations

Abbreviation	Explanation
CS	Cybersecurity
DPIA	Data Protection Impact Assessment
GA	Grant Agreement
GDPR	General Data Protection Regulation
GDPRCSA	GDPR Compliance Self-Assessment
IAM	Identity and Access Management
IdMS	Identity Management System
MISP	Malware Information Sharing Platform
MVP	Minimum Viable Product
NAS	Network Attached Storage
Observatory IE	Observatory Information Exchange
Observatory KB	Observatory Knowledge Base
OTM	Organisational and Technical Measure
PA	Processing Activity
PDP	Personal Data Protection
ROPA	Record Of Processing Activities
SAE	Self-Assessment Engine
SCORE	Security Capability-Oriented Requirements Engineering
UI	User Interface

Executive Summary

This report accompanies D4.1, a demonstration, which includes the services which are part of the **Minimum Viable Product** (MVP) of the SENTINEL project. It is important to note that the technologies and offerings, whose technical descriptions are included in D4.1, are integrally linked and mostly interdependent with their counterpart technologies in D2.1 and D3.1, all of which contribute to and participate in the integrated MVP version which is described in D5.4 (The SENTINEL Minimum Viable Product).

In more detail, the SENTINEL services taking part in D4.1 are: (a) The AIRBUS CyberRange simulations; (b) the SENTINEL DPIA framework and plugin, (c) The SME self-assessment engine for privacy and personal data protection which, in turn, comprises both (i) the SME profiling service; and (ii) the SME self-assessment service; and (d) the SENTINEL observatory. D4.1 is dedicated to technical descriptions of MVP of these services which are released in M12.

The technical work leading up to the SENTINEL Services (WP4) is driven from the project's baseline, defined in WP1, and brings together work carried out in both WP2 (The SENTINEL privacy and personal data protection technologies) and WP3 (The SENTINEL digital core). D4.1 attempts a concise presentation of the four aforementioned services, through a brief description of their purpose, role in the SENTINEL architecture, role in the SENTINEL MVP and any relevant technical details, including those with reference to design, implementation, deployment, and testing.

1 Introduction

1.1 Purpose of this document

1.1.1 Scope

The focus of this deliverable is on the demonstration of the SENTINEL Services which participate in the project's MVP demonstrated in M12. This report provides concise and summary technical descriptions of the different modules and services which are part of WP4 and, combined, form the MVP version of the SENTINEL Services.

The detailed MVP architecture, integration and use cases are outside the scope of the present report and are part of deliverable 'D5.4 – The SENTINEL Minimum Viable Product'. This report, however, provides technical descriptions of the SENTINEL components which participate in all WP4 (The SENTINEL services) tasks and provides a reference point for both D5.4 and the actual live demo of the MVP. Specifically, there are four (4) services deployed as part of this release: (1) The AIRBUS CyberRange simulations; (2) The SENTINEL DPIA framework and plugin; (3) The SME self-assessment engine for privacy and personal data protection which, in turn, comprises both a. the SME profiling service; and b. the SME self-assessment service; and (4) the SENTINEL observatory.

1.1.2 Contribution to WP4 and project objectives

D4.1 is the first written documentation of the work carried out in SENTINEL's WP4. Even though in the GA, only two out of four tasks (T4.1 and T4.3) had concrete technical deliverables to be included in the MVP release, it was eventually made possible for all tasks to deliver a minimum functional service, corresponding to each of the work package's four (4) tasks, as part of the MVP release.

The technical work presented in D4.1 is aligned to the Objectives of WP4, in the manner detailed below:

Objective 1. *The SENTINEL Cyber Range testbeds for simulations and training*

This objective is addressed through the provision of the CyberRange platform by SENTINEL consortium partner Airbus Cybersecurity (ACS). The Airbus CyberRange provides a simulation and testing environment (testbed) for an educational, collaborative, and hands-on training for real-world cybersecurity scenarios for SMEs, easily configurable to their specific on-premises or Cloud infrastructures. The Airbus CyberRange is initially deployed as an external service in the present SENTINEL MVP release, leveraging OAuth SSO user authentication integrated with the SENTINEL IAM (IdMS). The platform will be subsequently enriched with predefined scenarios and templates suitable for SMEs. This work is presented in Section 2 of this report.

Objective 2. *The SENTINEL data protection impact assessment (DPIA) framework*

This objective is addressed through the provision of the DPIA framework by SENTINEL consortium partner Sphynx Technology Solutions (STS). The STS DPIA framework is a questionnaire-based self-assessment plugin integrated with SENTINEL to allow SMEs to identify (via self-assessment) and minimise (via recommendations) the risks associated with the personal

data processing activities with which the SMEs have populated their SENTINEL profile. Although, the DPIA framework was not initially marked for inclusion in the SENTINEL Services MVP, it has been made possible for a basic functional version of this plugin to be designed, implemented and included in this release. This work is presented in Section 3 of this report.

Objective 3. *Tailor-made and intelligent requirements analyses, followed by the design and deployment of the necessary training sessions and a smart self-scoring mechanism*

This objective is addressed by consortium partner IDIR, through (a) the maturing of the SCORE (Security Capability-Oriented Requirements Engineering) framework and metamodel for Cybersecurity (CS) and Personal Data Protection (PDP), researched and developed in T1.1, into a flexible and automated tailored requirements approach for SME profiling utilising the notion of patterns; (b) the design and implementation of the SENTINEL Profile Service which is responsible for realising the common domain model and providing centralised storage and retrieval services for all data related to the participant organisations; and (c) the design and implementation of the SENTINEL Self-assessment Service which is responsible for assessing both (i) the eligibility status of the organisation and its processing activities for various self-assessments or recommendations, and (ii) the initial risk score of these personal data processing activities. This work is presented in Sec. 4 of this report.

Objective 4. *The SENTINEL Observatory and knowledge base*

This objective is addressed through the design and development, from the ground up, of the Observatory and Knowledge Base, by consortium partner ITML. For the MVP release, the initial version of the Observatory provides SENTINEL users with access to a centralised threat intelligence knowledge base for cybersecurity, privacy and personal data protection, aggregating and exchanging data in real-time among open security data platforms. This work is presented in Section 5 of this report.

1.1.3 Relation to other WPs and deliverables

The technical work in the fourth work package (WP4) is:

a) **Driven** by the project's baseline, defined in detail in 'WP1 – The SENTINEL baseline: Setting the Methodological Scene'; specifically detailed in 'D1.1 – The SENTINEL baseline' and 'D1.2 – The SENTINEL technical architecture' which, together define both the state of the art on which the project requirements are based and the refined architecture for the SENTINEL framework.
b) **Interrelated** to the work concurrently developed within work packages WP2 and WP3 of the project. This technical work is detailed in deliverables D2.1 and D3.1 respectively. Specifically, in WP2 and WP3,
    a. T2.1 "The SENTINEL data protection and data privacy compliance framework" provides the GDPR Compliance Self-Assessment plugin (GDPR CSA) which is the project's first self-assessment component, operating alongside the DPI Self-Assessment plugin (DPIA) developed in T4.2. The outputs of both modules are evaluated by SENTINEL's Self-Assessment Service and stored by the Profile Service, both parts on the project's architecture and developed under T4.3.
    b. The outputs of T2.3 "Contributed cybersecurity components" are initially considered in an attempt for T4.3 to design and store cybersecurity data in the

form of a cyber asset inventory which is to be associated with OTMs and specific measure implementations to be recommended by Core (WP3) to be part of the drafted policy.

c. The Recommendation Engine (T3.3) is exchanging data with the Profile Service (T4.3) in order to be informed and to inform the Policy Drafting Module (T3.4) which, in turn, also need to read data from the Profile Service (T4.3) when requesting inputs for how to draft policy and, also, write data back to it when saving this policy for later reference.

c) **Tightly linked** to the integration task of the project (T5.2 "Continuous integration towards the realisation of a complete system"). Following the work performed in T5.2 and described in detail in D5.4: "The SENTINEL Minimum Viable Product", all the technical components and services developed, not just within WP4 but also the tightly coupled technical tasks in WP2 and WP3, are integrated in a unified manner, a defined infrastructure and are able to support clearly defined use cases with the participation of end-users in a real-world demonstration.

Finally, D4.1 will provide the necessary foundation for the development of later versions of reports documenting the progress of the SENTINEL services, namely D4.2: "The SENTINEL services: Full-featured version" (due M18) and D4.3: "The SENTINEL services: Final product" (due M30).

1.2 Structure of the document

D4.1 follows a structure closely aligned with the workplan and task structure of WP4. Four distinct sections (Section 2-Section 5) directly correspond to the technical work completed in T4.1, T4.2, T4.3 and T4.4, respectively, with section 6 dedicated to overall summary of this report. Specifically:

- *Section 2* describes the AIRBUS CyberRange simulations.
- *Section 3* describes the SENTINEL DPIA framework and plugin.
- *Section 4* describes the SME Profiling Methodology and the SENTINEL Self-Assessment & Profile services.
- *Section 5* describes the SENTINEL Observatory.
- *Section 6* summarizes the report and recommends future action.

1.3 Intended readership

This report (D4.1) is released publicly and linked to the demonstration of the project's MVP release in month 12 (M12). It is thus considered an accompaniment to D5.4: "The SENTINEL Minimum Viable Product" which holds a full technical description of the integration activities and the use cases supported by the MVP demonstration. D4.1 is therefore addressed to

- Consortium technical partners looking for reference to technical components developed within WP4.
- Internal (consortium) and external reviewers.
- End users looking to understand the purpose, role and technical details of the SENTINEL services.
- The general public, including other research projects and activities.

2 Advanced CyberRange simulations

SENTINEL's Cyber Range simulation environment (testbed) is based on the CyberRange platform provided by consortium partner **Airbus Cybersecurity** (ACS).

2.1 The CyberRange testbed

The CyberRange is a simulation platform that can be used either for testing systems before on-site integration, optimizing cyber-defence strategies or training the end-users.

2.1.1 CyberRange overview

The platform offers an existing library of virtual machine and docker, to make it easier to start modelling SME's IT infrastructures to be simulated. There is also the possibility to integrate an external virtual machine or docker and connect physical equipment to the virtual network.

The CyberRange can be deployed for large infrastructures, but it can also be used by smaller enterprises without access to CS experts. Thanks to the user interface, it is easy, even for non-expert IT staff to replicate and deploy the SME's infrastructure in a simulation. With a drag and drop interface, the user is able to deploy predefined workstation and network templates. This offers the possibility for SMEs, to enable self-assessment and discover vulnerabilities.

The CyberRange is composed of *Workzones*. A Workzone is a set of resources (memory, CPU, network). All work zones are isolated from each other and give the possibility to efficiently deploy networks and hosts. The CyberRange can also be accessed remotely. From the web interface, users can access different Workzones and open remote console to visualise the virtual machine and interact with it. For example, a trainer can launch and manage a cyber scenario for trainees in real-time. The Figure 1 below show an example of a Workzone, with a topology deployed.

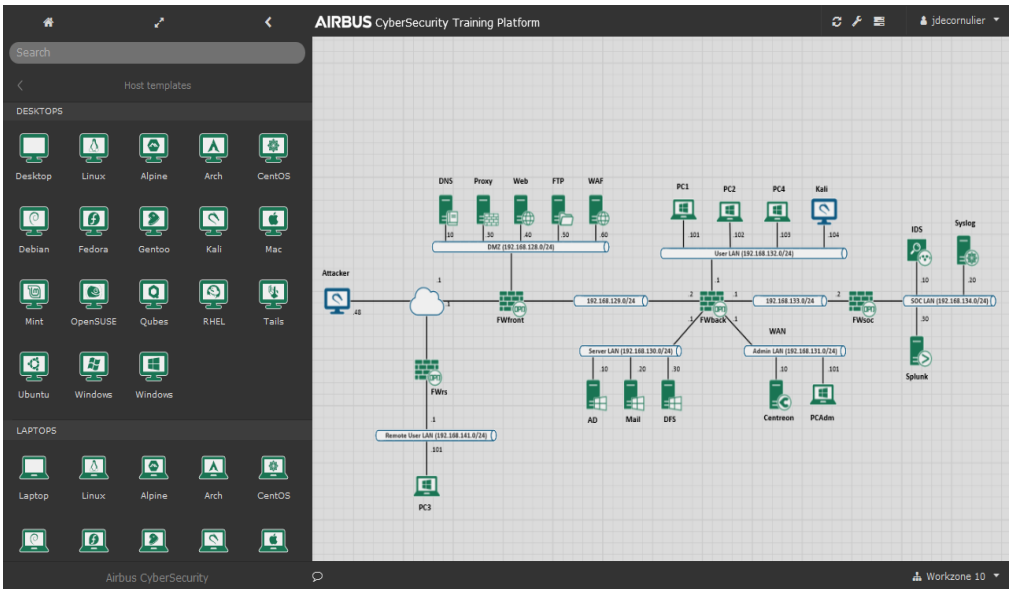

Figure 1. A CyberRange Workzone

2.1.2 CyberRange technical specifications

CyberRange is a platform composed of physical servers and switches, hosting VMware vSphere Infrastructure. The physical infrastructure of the CyberRange platform is located at the Airbus CyberSecurity premises (Elancourt, France). The CyberRange platform is mainly composed of one switch (CyberRange CR16), one NAS (Network Attached Storage) and several servers that host the virtual platform. Network access to the infrastructure is protected by a firewall which allows connecting other systems from different rooms of Airbus premises or even from the Internet so that SENTINEL participants can access the virtual platform.

2.1.3 CyberRange integration with SENTINEL

The CyberRange platform provides an OpenID plugin to authenticate users against SENTINEL's IdMS (identity and access management system and SSO) based on Keycloak. From the SENTINEL interface, users may utilise a web link to redirect to the CyberRange dashboard and are seamlessly authenticated via OpenID.

The CyberRange platform exposes a public page that acts as an OpenID client. This page accepts an "authorization_code" and is configured to call the Sentinel OpenID/SSO provider.

When a SENTINEL user clicks-through to the Cyber Range link in MySentinel, the following workflow is enabled:

- SENTINEL redirects to the CyberRange login page.
- SENTINEL **/authorize** endpoint creates an **authorization_code** and redirects the user to the CyberRange **public page** (API calls).
- CyberRange reads the **authorization _code** and sends it along with its client configuration to SENTINEL **/oauth/token** endpoint.
- SENTINEL responds with an **ID Token** (JWT) and an **Access Token** (JWT). The **ID Token** contains all required user information.
- If the user does not exist, the CyberRange inspects it and puts it in the right groups. The groups allow to give permissions in the platform to the user. Each user will be able to access only their dedicated workzone. To prevent overlap with groups from other projects, the name of all user groups used by SENTINEL will start with "sentinel_" pattern in the CyberRange.

2.1.4 CyberRange functionality

Open remote access is provided to the SENTINEL consortium members and participant SMEs so that they can replicate their infrastructure, test their security systems, and simulate attack scenarios.

Figure 2 below shows the graphical user interface of the virtual platform:

- On the middle: Workbench for visualization and interaction with the current Workzone.
- On the left part: Navigation Menu to access the library.

Figure 2. CyberRange Graphical Interface

A network infrastructure library is available to deploy virtual machines and containers into the Workzone:

- Applications Templates: This category contains applications, mainly containers.
- Networks: This category is used for basic network solutions like routers or firewalls; it can be either containers or virtual machines.
- Operating Systems: This category contains basic Operating Systems (Debian, Centos, Ubuntu, Kali Linux, Windows …); it can be either containers or virtual machines.
- Topologies: This category contains topologies with multiple hosts (containers, virtual machines) and networks.

The deployment of Operating Systems is carried out by a simple drag and drop into the Workzone. Then the user must fill some parameters such as hostname, CPU cores, memory size and description to create the host.

As for the deployment of networks, it is also performed via a drag and drop into the Workzone. Then the user can fill some parameters such as IP address, network mask, gateway and description to create the network.

With the CyberRange platform, it is possible for a user to execute commands in hosts:

- For containers: With a double-click on a container, or right-click and Execute command, the user can then type a command in the command text box and click on the Execute button.
- For virtual machines: The user can open a remote console with a double-click on it. The console screen is shared between all the users of the *Workzone*.

Once a host and a network are deployed into the Workzone, a user can connect them together by simply clicking on the host and then on the network. The user must then fill some parameters such as IP address and gateway to create the network adapter.

It is also possible to connect two networks: This requires deploying a Router from the Network section in the library. To connect both networks, a user must connect the Router to the first network and then to the second network. The user must then fill some parameters such as IP address and gateway to create the network adapter.

A topology created can be saved; a user must select the networks, routers and hosts items to save the topology. Finally, the user enters a name corresponding to the topology. The CyberRange platform allows to redeploy a topology saved by dragging and dropping it into the Workzone (The topology is available in the section Topologies of Network Infrastructure).

2.2 CyberRange demo scenarios for SME awareness & training

The CyberRange platform will provide an educational, collaborative, and hands-on cyber platform for simulating SME real-life cybersecurity scenarios. Based on the result of the self-assessments conducted by the SME, simulation and training using CyberRange will be recommended by SENTINEL.

SMEs will have full access to all CyberRange platform features. They will have access to a library of attacks, life generators and scenarios that is localized in the navigation menu, in the section Actions and Scenarios. For example, an SME user may configure an attack, such as a port scan as illustrated in Figure 3, by selecting “Attacker” and “Target” parameters, and then start the attack.

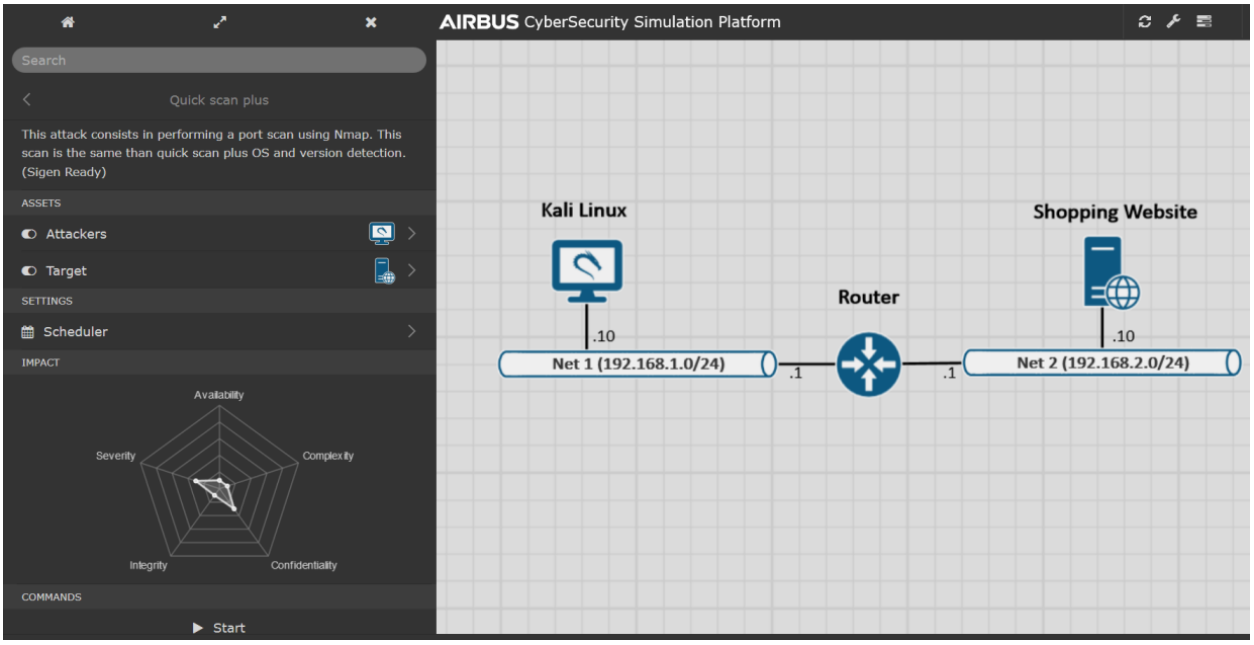

Figure 3. Quick scan plus attack

Generating life traffic is a functionality of the CyberRange. Generators, such as HTTP Generator, are available on the library in Life Traffic category of Actions and Scenarios section. A user must fill the fields Source and Destination to run a life traffic generator.

Attacks and life traffic can be combined into a scenario. Scenarios can be configured and are available on the Actions and Scenarios menu. As shown in Figure 4, the scenario graph contains several actions, either attacks or life traffic, represented by squares. The actions results are accessible by clicking on them after the start of the scenario.

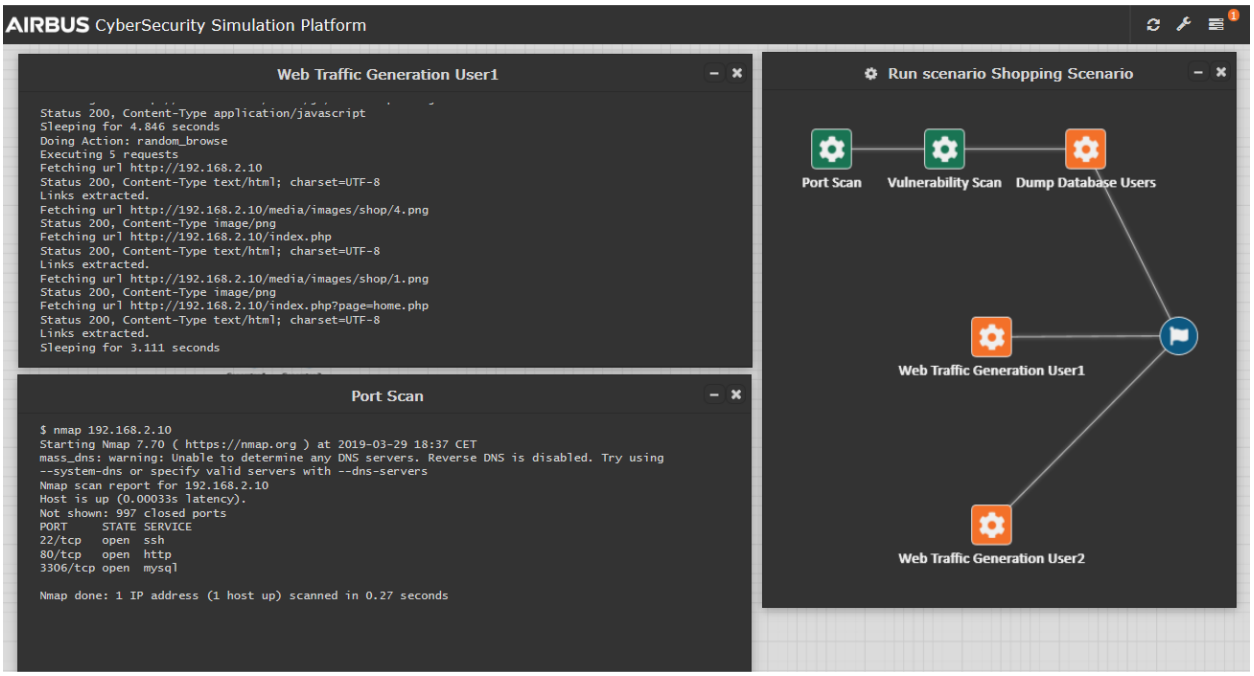

Figure 4. Example of CyberRange scenario

SME infrastructure will be replicated, and realistic scenarios subsequently created, based on potential vulnerabilities, to provide as many simulation and training templates as possible, on which SMEs may operate the CyberRange requiring just a basic IT security understanding, and not advanced cybersecurity expertise.

3 The SENTINEL Data Protection Impact Assessment

3.1 Overview

SENTINEL's Data Protection Impact Assessment (DPIA) toolkit was designed to allow SMEs to identify – through assessment - and minimise – through recommendations – the risks associated with their personal data processing activities. Although, based on the DoA, the DPIA was not explicitly marked for inclusion in the MVP, it has been made possible for a basic functional version of this plugin to be designed, implemented and included in the release.

The MVP version of the DPIA toolkit is based on state-of-the-art tools[1] and questionnaires[2] tailored to the needs of SENTINEL. Nevertheless, SENTINEL's DPIA Toolkit is responsible for constructing the DPIA questionnaire and subsequently, for calculating the Processing Activities' risk based on the participant's responses. The MVP questionnaire includes 19 questions, where each question may have one or more (1..*) options. Each option has a specified impact and likelihood. After an end-user submits the questionnaire, the DPIA Toolkit calculates the risk (based on the likelihood and impact of the selected options per question), and provides some qualitative, and quantitative metadata which can be used both for the presentation and storage of the self-assessment results and for the subsequent recommendations.

3.2 Architecture and technical specifications

The MVP version of the DPIA, consists of two components, (a) the DPIA toolkit, and (b) the DPIA database (see Figure 5). The former generates the DPIA questionnaire, get the DPIA response, and calculate the risk, while the latter stores the questionnaire and the results of the DPIA process.

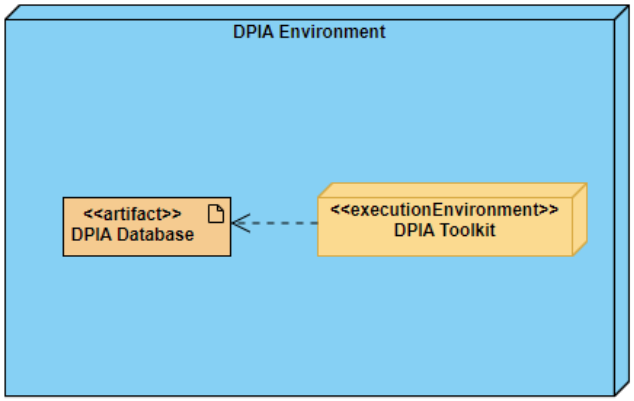

Figure 5. DPIA Environment (Deployment Diagram)

For the realization of the MVP version of the DPIA Toolkit, a sequence diagram (see Figure 6) showing the interactions between an end-user and SENTINEL's modules (incl. the DPIA Toolkit) was created.

[1] https://www.cnil.fr/en/privacy-impact-assessment-pia

[2] https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-impact-assessments/

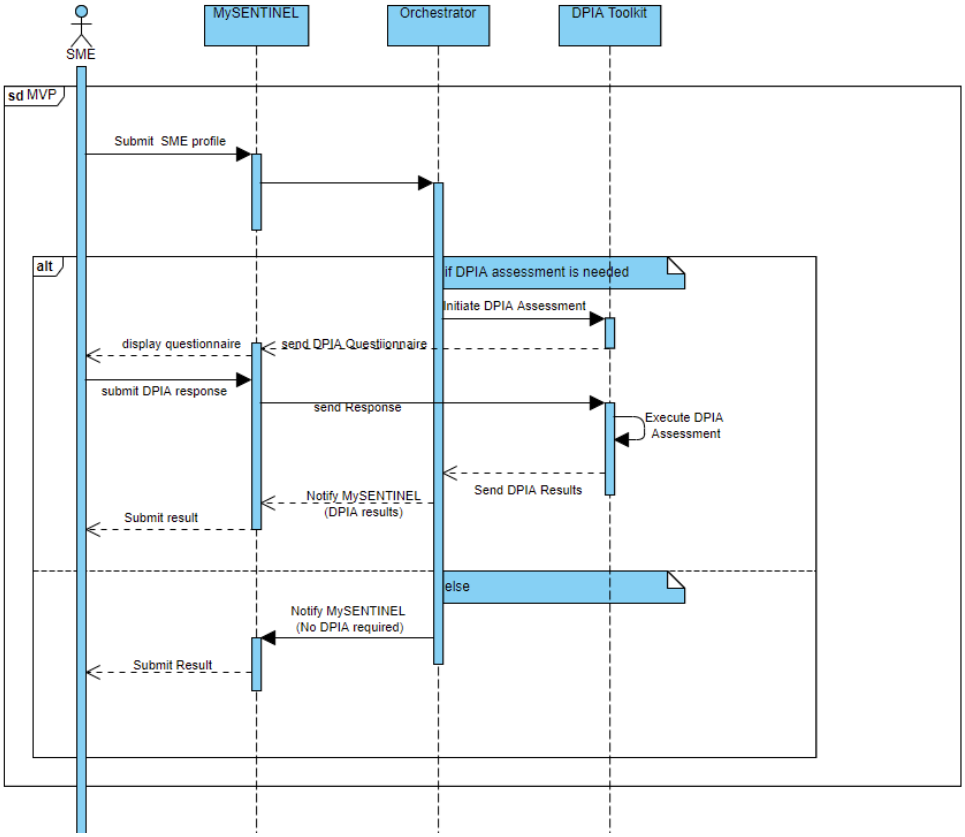

Figure 6. DPIA sequence

For the implementation of the MVP for the DPIA, the following technologies were used:

- **DPIA Toolkit**: This module is implemented in Java, using the Spring Boot framework[3]. The toolkit is also responsible to serve several RESTful APIs that are responsible for submitting (POST) the DPIA questionnaire, retrieve the responses to the questionnaire (GET), and retrieve the DPIA Results (GET). To describe the APIs, OpenAPI v3[4] was utilized.
- **DPIA Database**: The database response to store the DPIA questionnaire and results. The technology used for the database is Postgres[5], a powerful, open-source object-relational database. Part of the DPIA results is also stored in the SENTINEL Profile Service.

Lastly, the DPIA Toolkit was containerized using Dockerfile[6], and docker-compose[7].

[3] https://spring.io/projects/spring-boot
[4] https://swagger.io/specification/
[5] https://www.postgresql.org/
[6] https://docs.docker.com/engine/reference/builder/
[7] https://docs.docker.com/compose/

4 The SENTINEL Profiling and Self-Assessment Services for privacy and personal data protection

4.1 The SENTINEL SME profiling model and methodology

4.1.1 From SCORE to SME profiling: The SENTINEL profiling meta-model

In D1.1 the methodology known as SCORE (Security Capability Oriented Requirements Engineering) has been defined as the conceptual baseline for specifying CS and PDP requirements. The conceptual foundation of SCORE was further elaborated in WP4 to incorporate profiling requirements, resulting in an augmented metamodel, as shown in Figure 7.

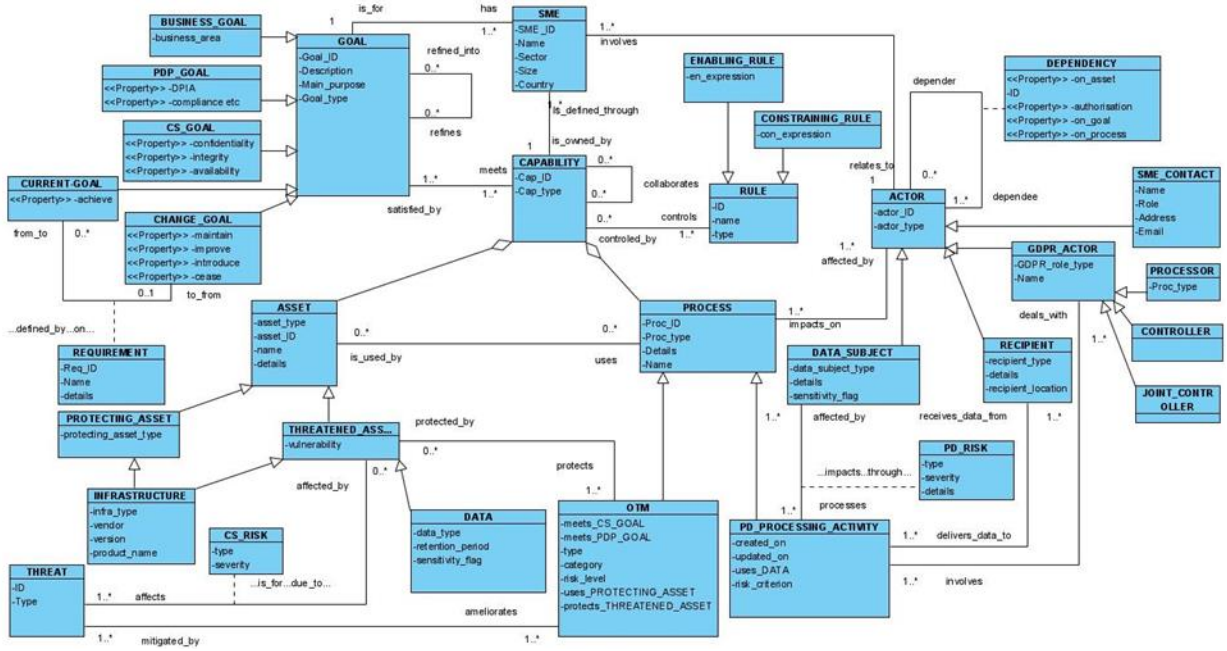

Figure 7. The SENTINEL profiling metamodel

In terms of CS and PDP requirements the metamodel helps to record perceived THREATS that are identified by business users as having an IMPACT on CURRENT CAPABILITIES. Analysis of such threats and their potential impact will lead to the definition of new business goals (CHANGE GOALS) and their corresponding DESIRED CAPABILITIES leading to THREAT MITIGATION.

As shown in Figure 7, a CAPABILITY is defined as an aggregation of PROCESSES using ASSETS. Threat mitigation is based on the implementation of appropriate organizational and technical measures (OTMs) pertaining to a DESIRED CAPABILITY aiming to protect the ASSETS being affected. For example, data breach is a THREAT that has a high impact on the organisation's capability to process personal data. Strengthening the current goal of secure data processing is of improve type REQUIREMENT that should be met by the desired personal data processing capability, which improves the current capability (CAPABILITY TRANSFORMATION)

by implementing a number of OTMs such as ‘enforcing an access control policy’, ‘authentication and access control’, etc., thus mitigating the THREAT.

4.1.2 Towards tailor-made requirements analyses

The conceptual framework detailed in section 4.1.1 offers a common terminology for capturing risk associated with the processing of personal data and for identifying the required CS and PDP capabilities, during profiling. Using this knowledge, we consider the use of a ***pattern-driven approach*** for tailor-made requirements analysis, whereby elicited knowledge is used to recommend appropriate OTMs and other resources (i.e., trainings and plugins).

Patterns as a means to encapsulate and communicate proven security and privacy solutions, is an active and growing field of research. In SENTINEL, patterns are used as a means to assist SMEs to identify the appropriate OTMs that ought to be present in their CS and PDP policy. Patterns are described in terms of the relevant concepts defined at a conceptual level in SCORE, as shown in the pattern’s conceptual model in Figure 8.

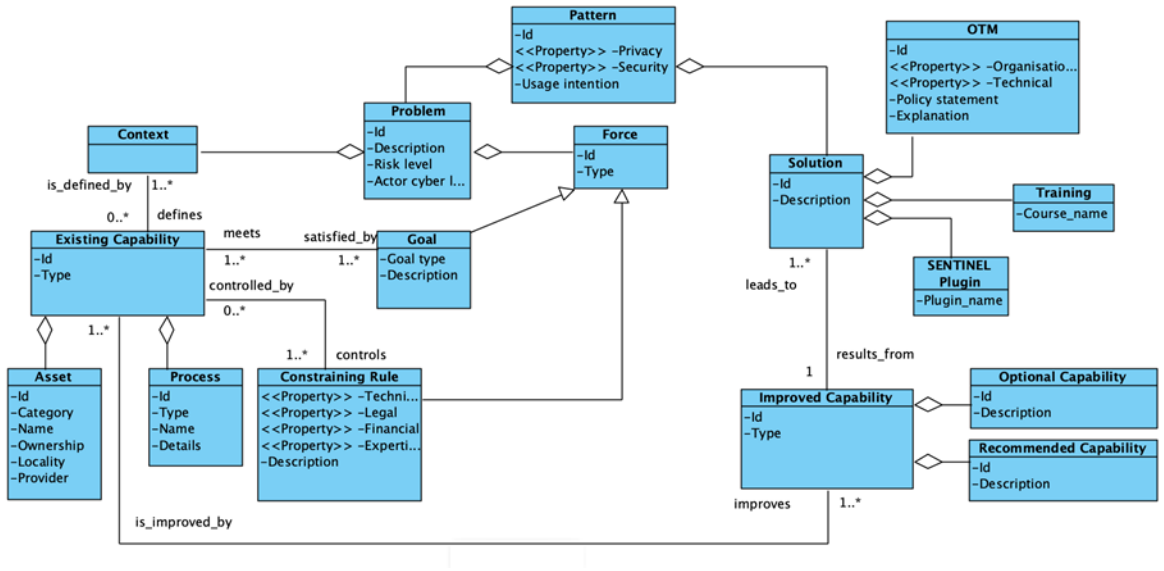

Figure 8. Definition of a pattern in SENTINEL

To demonstrate the use of the template we have used the information from the TIG pilot case which was reported in D1.1. In the TIG pilot case, an existing capability was defined as “Service Providing”, which is further specified in terms of the “Data sharing with social agencies” process, which uses two assets namely those of “Service User Data” and “Cloud Information System”.

A pattern is defined in terms of two key aspects namely those of Problem and Solution. A problem is a confluence of Contexts and Forces. A Context defines the situation in which a problem occurs for example, an SME having the capability of ‘service providing’. A Force defines those issues that influence the problem, and which must be resolved, for example, meeting the goal of “improving the protection of service user data”, whilst “conforming to government regulations".

```
If there exists a Context defined by
 an Existing Capability Ci
   involving Assets A1, .. An
   AND Processing Activities PA1, ..,PAn
 with Problem of
   Risk Level RLi
   AND Actor Cyber Level ACLi identified by
    a set of Forces
     involving Goals G1, G2, ..Gn
     AND/OR Constraining Rules R1, R2, ..,Rn
Then apply Solution Si
 involving OTMi AND/OR TRAININGi AND/OR PLUGINi that
    improve Goals G1, G2, ..Gn
    AND/OR meet Constraining Rules R1, R2, ..,Rn
 leading to Recommended Capability RCi
AND Optional Capabilities OC1, OC2, ...OCn
```

Figure 9. A generic rule-based description of the pattern template

A Solution is made up of three elements namely those of policies (OTM), awareness practices (Training) and technology components (Software Plugin). Applying the Solution would lead to some Improved Capability, defined in terms of Recommended Capability and Optional Capability.

Finally, the recommended OTMs forming the solution include “To enforce access control policy” and “To provide third-party delivered and monitored CS services”.

Considering the TIG pilot case, we can define an instance of the pattern template as shown in Figure 10.

```
If there exists a Context defined by
 an Existing Capability Service Providing
  involving Assets Cloud IS (sw_saas, cloud, no_owned_asets, google)
  AND Processing Activities Data exchange with social agencies
 with Problem of Risk Level risk high
  AND Actor Cyber Level intermediate identified by
   a set of Forces
   involving Goals Confidentiality of service user data
   AND/OR Constraining Rules CQC/CIW Regulations
Then apply Solution Si
 involving O1.H.1 (Semester PDP Policy Review Process)
  that  meet Constraining Rules CQC/CIW Regulations
  AND improve Goal Confidentiality of service user data
 leading  to  Recommended  Capability  O1  (org_policy_drafting_enforcing,
Defining and enforcing a policy)
  AND Optional Capability s_cloud_security (To provide third-party (Cloud)-
        delivered and monitored CS services)
```

Figure 10. Using the pattern template in the TIG pilot case

4.2 The SENTINEL Profile Service

The Profile Service plays a central role in the SENTINEL technical architecture, by providing centralised storage and dissemination services for data related to the participant organisations (SMEs).

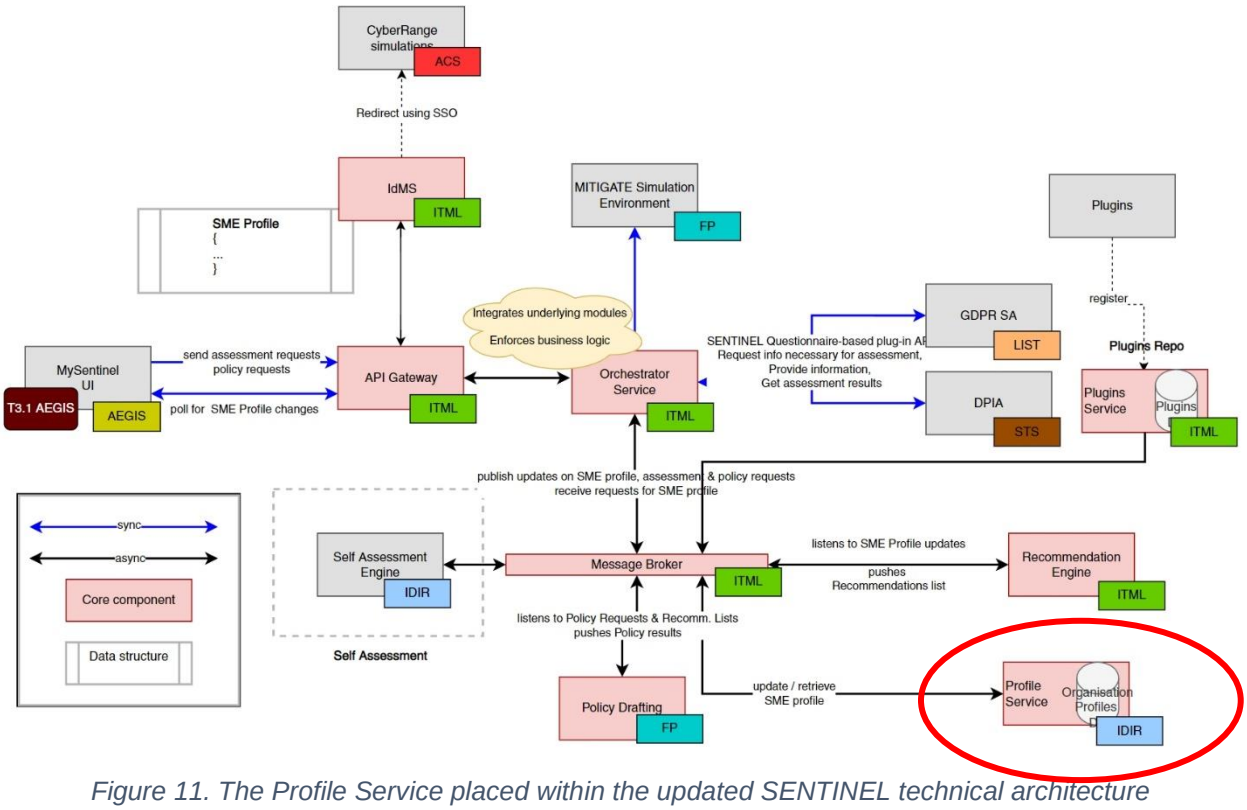

Figure 11. The Profile Service placed within the updated SENTINEL technical architecture

Specifically, the Profile Service is responsible for

(a) dynamically providing the definitions of the data required for the front-end (MySentinel) to populate the SME profiles, also described in the technical documents as 'questionnaires'. These data definitions pertain to (i) the SME organisation and, specifically to the core organisation data, the contact persons for personal data protection and the overall organisation asset profiling and (ii) the organisation's processing activities

(b) Implementing the SENTINEL domain model (Figure 12) for SME organisations and providing persistence for storing and fetching:
- Core organisation data
- Organisation contact persons
- Asset-related data (asset profiling)
- PD processing activities
- Internal PA-related data such as
  - PA risk
  - self-assessment and recommendations eligibility
- ROPA-related data (a permanent storage for processing activities)
- Self-assessment results

        - The output of the GDPR compliance self-assessment plugin (GDPRCSA)
        - The output of the DPIA self-assessment plugin
    - The output of the Recommendation Engine (Recommendations list)
    - The output of the Policy Drafting module (Policy results)

It should be noted that the data model implemented for the MVP instantiates part of the SENTINEL profiling metamodel proposed in Section 4.1.1. The full featured (M18) and final (M30) versions of the project will transition towards a more inclusive profiling approach which will address business goals and requirements coupled with CS and PDP ones and a direct mapping between cyber assets (inventory) and CS / PDP capabilities.
In order to describe the necessary APIs or the above, we used OpenAPI v3[8]. Appropriate Service endpoints enable SENTINEL to (a) Create Organisation; (b) Update Organisation; (c) Retrieve Organisation; (d) Create Processing Activity; (e) Update Processing Activity; (f) Retrieve Processing Activity; (g) Store Assessment Eligibility Results; (h) Retrieve Assessment Eligibility Results; (i) Store DPIA or GDPRCSA; (j) Retrieve DPIA or GDPRCSA; (k) Store Recommendation Results; (l) Retrieve Recommendation Results; (m) Store Policy Drafting Results; (n) Retrieve Policy Drafting Results; and (o) Provide the definition of fields for profile data capturing.

The SENTINEL Profile Service has been implemented as a microservice with Java 11, using Spring Boot[9]. It is also leveraging the SENTINEL Async API which uses RabbitMQ[10] as message broker. MongoDB[11] is used for the persistence of the data.

[8] https://swagger.io/specification/
[9] https://spring.io/projects/spring-boot
[10] https://www.rabbitmq.com/
[11] https://www.mongodb.com/

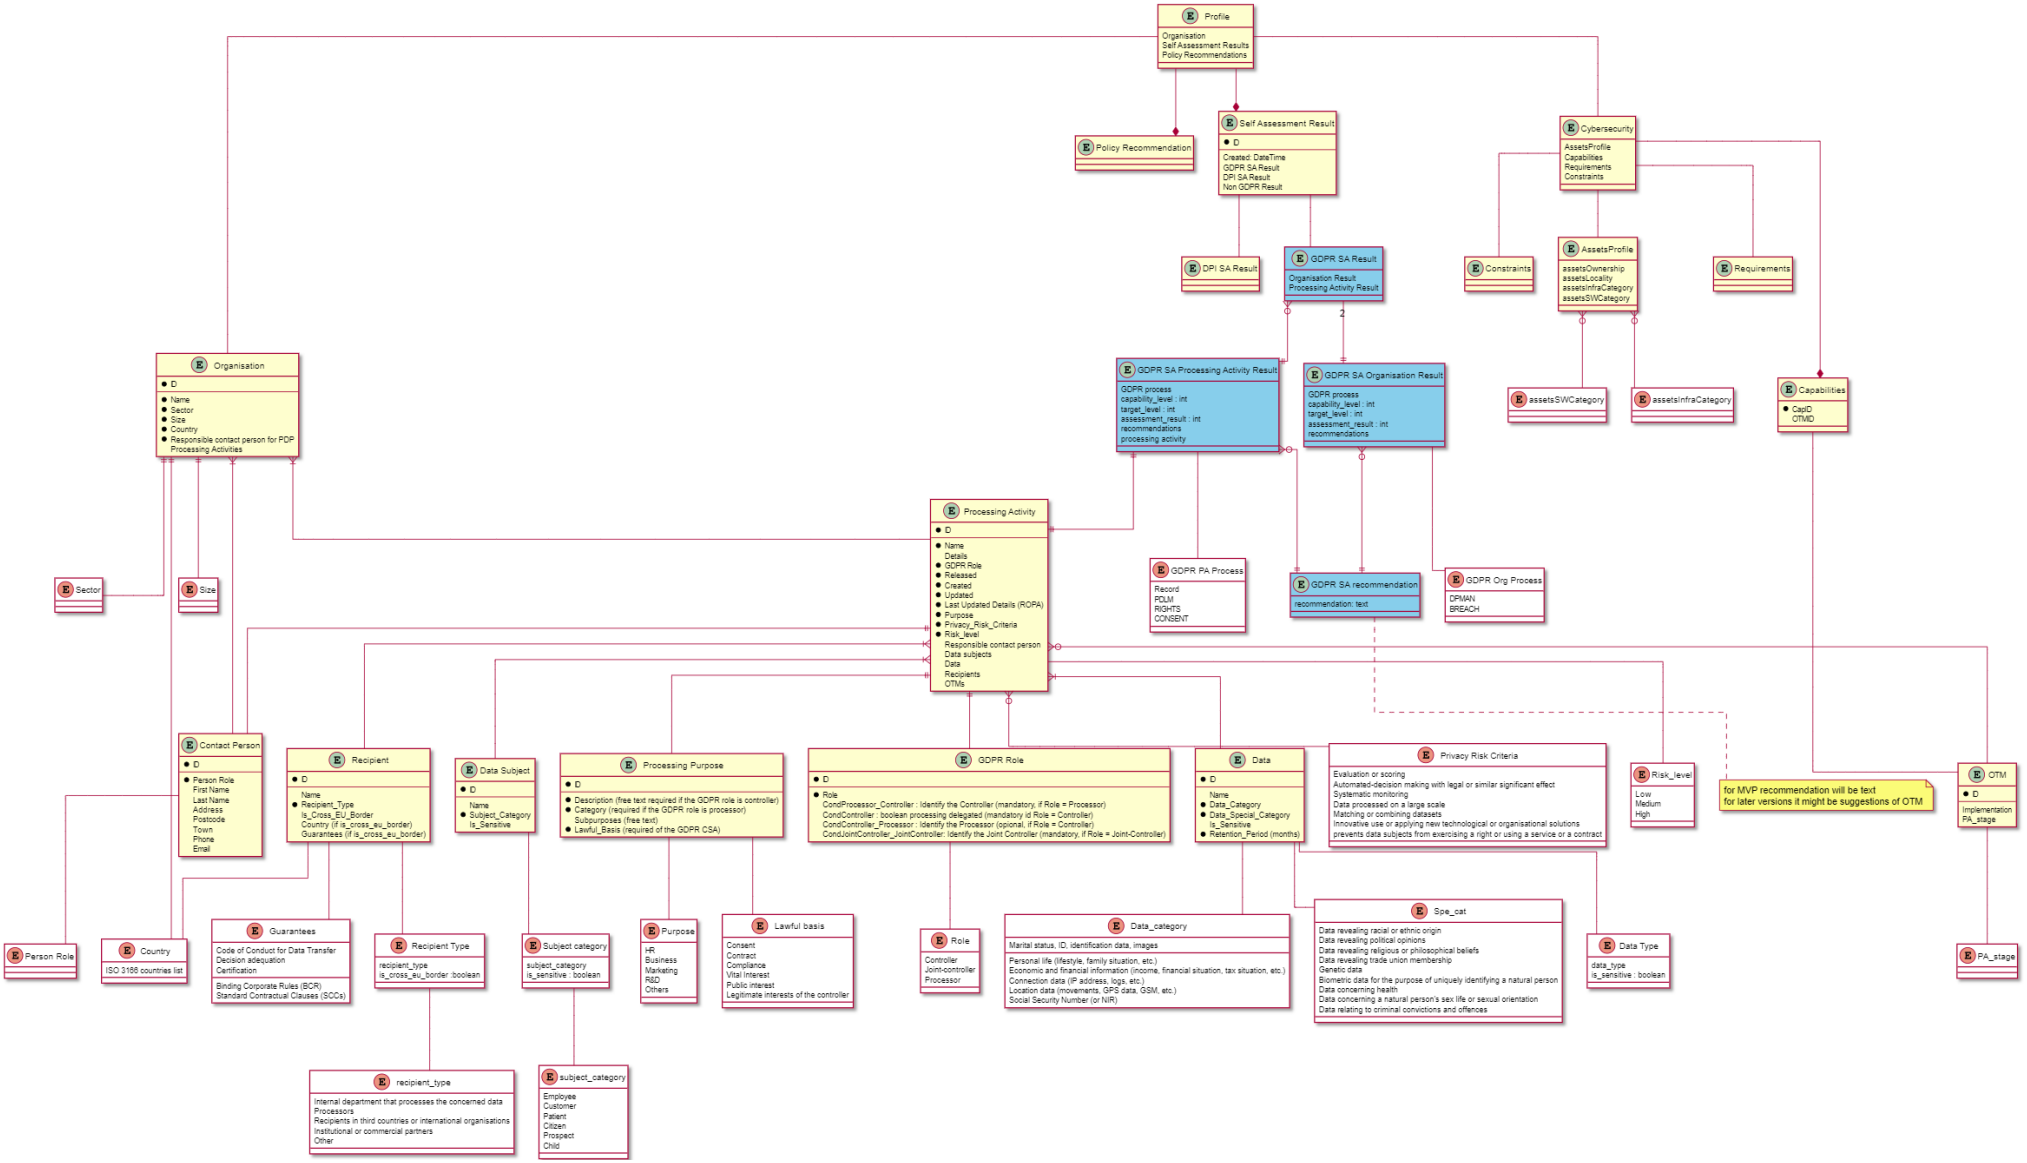

Figure 12. UML diagram of the common SENTINEL domain model supported at the MVP

4.3 The SENTINEL Self-Assessment Service

The Self-Assessment Service, also referred to in the architecture as the Self-Assessment Engine (SAE), is invoked every time the organisation profile is updated and is responsible for enabling specific SENTINEL assessment and recommendation workflows depending on the eligibility status of the organisation and its processing activities, and for performing an initial risk assessment.

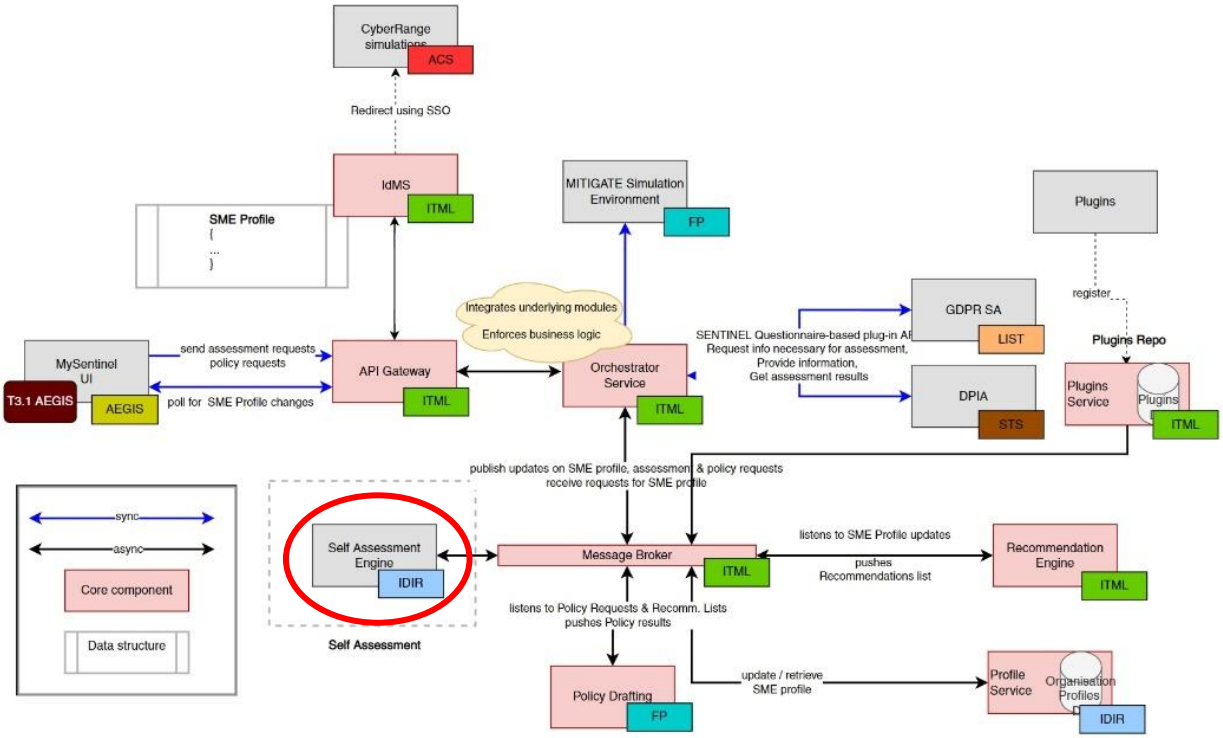

Figure 13. The SA Service placed within the updated SENTINEL technical architecture

In more detail, regarding its **eligibility assessment responsibilities**, the SAE decides:

- whether a processing activity is eligible for either self-assessment plugin (the GDPR Compliance Self-Assessment (GDPR CSA) and/or the DPIA Self-Assessment), by checking the updated processing activity for completeness. In practice, every successfully permanently saved processing activity (not saved as 'draft') is validated to be complete.
- whether the organisation is eligible for the cybersecurity assessment (CS self-assessment), by checking the organisation profile (core data and asset inventory) for completeness. *This functionality will be developed fully post-MVP* when SENTINEL will support robust cyber asset inventory capturing.
- whether the organisation is eligible for the policy recommendations workflow, according to the current state of the organisation profile and completeness of at least one PD processing activity.

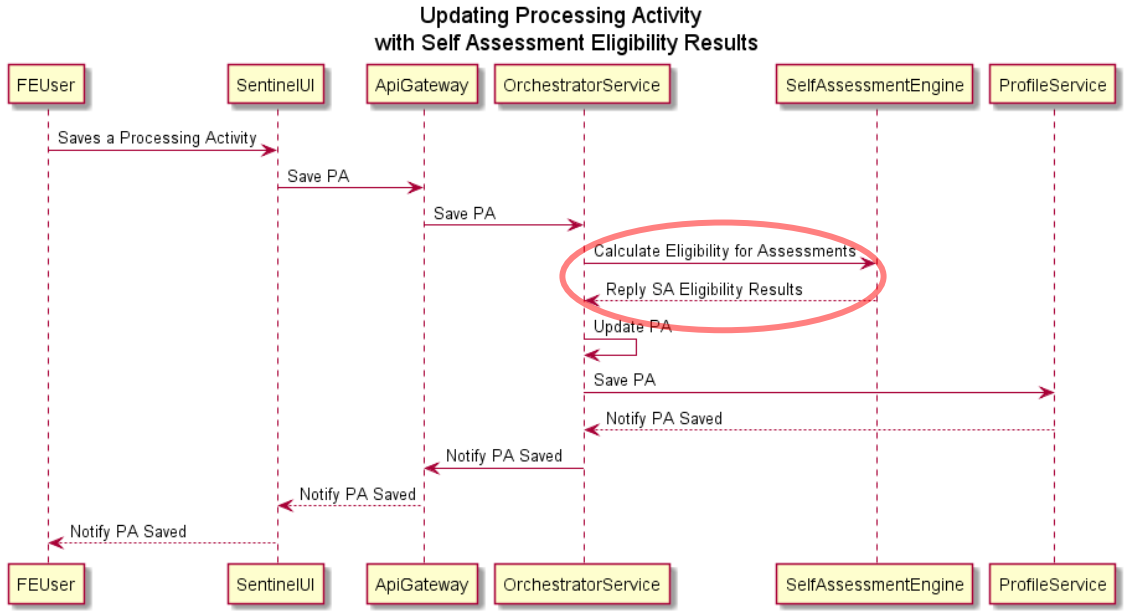

Figure 14. The system-wide sequence for checking the SA eligibility of PAs in SENTINEL

Regarding its **risk assessment responsibilities**, **the SAE initially calculates the risk level** (formerly referred to in the GA as the "RASE" score) of each "complete" Processing Activity by considering its privacy risk criteria, following a simple algorithm. The risker PA then lends its risk level to the organisation. By saying '*initially*' we mean that, because a DPIA assessment (which results in a more rigorous and granular estimation of the risk of each processing activity) has not been triggered yet, the system would benefit from an *initial* assessment of the *potential* risk involved in the organisation's personal data processing activities in order to enable a recommendations workflow. This assessment can be estimated by considering these processing activities' privacy risk criteria. The process followed for this assessment is as follows:

```
IF DPIA for this PA has NOT been invoked {
      PA_risk=0;
      IF (dataSubjects/subjectsCategory IS one of "Employees", "Patients",
      "Children") THEN PA_risk++;
      IF (data/dataSpecialCategories HAS at least 1 option checked) THEN
      PA_risk++;
      FOR EACH (privacyRiskCriteria CHECKED) PA_risk++;
      IF PA_risk>=2 THEN riskLevel="HIGH" ELSE riskLevel="LOW";
}
```

Figure 15. A simple algorithm for initially establishing the potential risk involved in each PA

The SAE has been implemented as a microservice with Java 11, using Spring Boot[12]. It is also leveraging the SENTINEL Async API which uses RabbitMQ[13] as message broker.

[12] https://spring.io/projects/spring-boot

[13] https://www.rabbitmq.com/

5 The SENTINEL Observatory

Although, based on the GA, the Observatory was not explicitly marked for inclusion in the MVP, it has been made possible for a basic functional version of this plugin to be designed, implemented and included in the release.

5.1 Overview

SENTINEL's Observatory is an intelligence knowledge hub, designed to provide three key functions:

- a centralised threat intelligence knowledge base for cybersecurity, privacy and personal data protection, exchanging data in real-time among open security data platforms as well as aggregating anonymised information collected or produced within SENTINEL, complete with a searchable KB, FAQ and collaboration tools.
- an open API platform to exchange threat intelligence between SMEs/MEs and SME associations.
- the potential reuse of policy elements when drafting new security and privacy policies for SME participants.

For the MVP release of SENTINEL, functions a) and b) are delivered, namely a) a first version of the centralized **Observatory Knowledge Base (Observatory KB)** and b) the **Observatory Information Exchange (Observatory IE)** module to exchange security related information between SENTINEL and external open security data sharing platforms.

The MVP version of the Observatory is accompanied by a user Interface provided by MySentinel (developed Task 5.1) and is closely coupled with to the exploration and monitoring of open cybersecurity data platforms (implemented in Task 3.1).

5.2 Architecture and technologies

The use case relevant to the Observatory is defined in *'D1.2 – The SENTINEL technical architecture'*, as **Use Case 06: Consulting the Observatory Knowledge Base**. The purpose of the use case is to provide registered SENTINEL SME/ME representatives to browse the SENTINEL Observatory Knowledge Base (Observatory KB) and access information about recently identified data and privacy breaches. The Knowledge Base is continuously updated and synchronised with external resources.

For the implementation of this use case, we embarked from the definition of a UML sequence diagram that shows the interactions between the end-user and SENTINEL's submodules, as depicted in Figure 16. The top part of the diagram shows the continuous updating of information stored in the Observatory KB, as collected from external platforms, in this case MISP[14], HELK[15] and the NIST vulnerability database[16]. This task is realized by the Observatory IE. The bottom

[14] https://www.misp-project.org/
[15] https://securitydatasets.com/consume/helk.html
[16] https://nvd.nist.gov/

part of the diagram shows the interactions between the end-user and the Observatory KB, through MySentinel's user interface that provide searching, filtering and browsing capabilities to help users detect entries of interest to their domain. Finally, the diagram shows interactions with API, a SENTINEL's infrastructure module that is described in *D5.4 – 'The SENTINEL Minimum Viable Product'*, where the execution and showcase of this use case can be found.

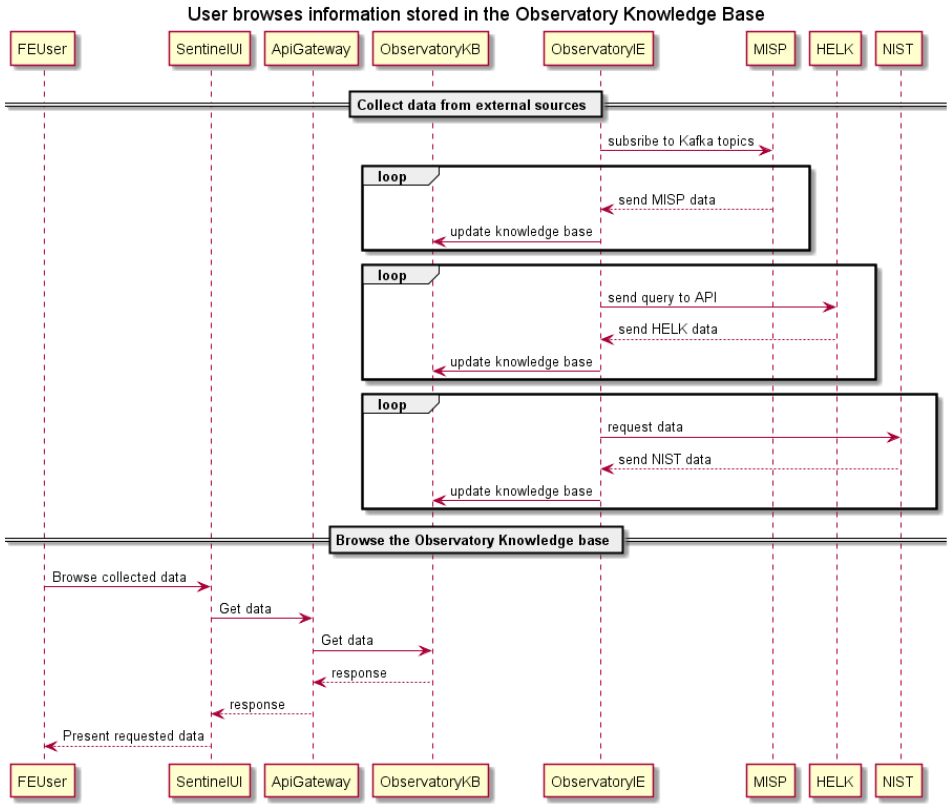

Figure 16. UML sequence diagram for the Observatory use case

For the implementation of the participating components, the following technologies have been employed:

- Observatory KB: An Elasticsearch[17] instance has been selected as it is both flexible for storing different types of documents, and performant in terms of searching and filtering over large amounts of stored data.
- Observatory IE: This module is implemented as a Java 11 service that implements queries to the APIs offered by the external data source mentioned above. For the MVP version, the MISP API client has been fully developed.
- Sentinel UI: The UI for the Observatory has been implemented using React[18].

[17] https://www.elastic.co/
[18] https://reactjs.org/

6 Conclusions

D4.1 summarises the technical work towards delivering the services in the context of WP4 of SENTINEL. All four tasks in WP4 correspond to a number of either user-facing services or internal modules implementing core parts of the SENTINEL architecture.

The list below summarises these technical deliverables for each of the four (4) tasks which, in turn, correspond to sections 2, 3, 4 and 5 of D4.1:

- WP4 : The SENTINEL services
    - Task 4.1: Advanced CyberRange simulations and training for SMEs/MEs
        - The Airbus CyberRange
    - Task 4.2: Data protection Impact assessment and assurance
        - The STS DPIA self-assessment framework (DPIA toolkit and DPIA database)
    - Task 4.3: Self-assessment and RASE scoring engine
        - The SENTINEL SCORE-based SME profiling model and methodology
        - The SENTINEL Profiling Service
        - The SENTINEL Self-assessment Service
    - Task 4.4: The SENTINEL Observatory
        - The SENTINEL Observatory

The presentation of each of the items above, together with the rest of the technical deliverables D2.1 and D3.1 developed and released concurrently with D4.1, serve as technical reference for SENTINEL's MVP release due M12 of the project. The detailed technical documentation of the integration tasks towards the MVP release is provided in D5.4: "The SENTINEL Minimum Viable Product".

Going forward, D4.1 will transition into its two future versions, namely D4.2: "The SENTINEL services: Full-featured version", due M18 and D4.3: "The SENTINEL services: Final product", due M30 which, in a similar fashion, will correspond to the future, more mature version of the WP4 technical deliverables, integrated in the interim and final versions of SENTINEL respectively.