Horizon 2020 Programme

H2020-SU-TDS-02-2018
Trusted digital solutions and Cybersecurity in Health and Care



DATA-PROTECTION TOOLKIT REDUCING RISKS IN HOSPITALS AND CARE CENTERS

Project Nº 826284

ProTego

D7.4 Report on the final results of experiments and demonstrations

Responsible:	Inetum (GFI INFORMÁTICA before January 2021)
Contributors:	All partners with the leadership of Inetum
Document Reference:	D7.4 Report on the final results of experiments and demonstrations
Dissemination Level:	Public
Version:	Version 1.0
Date:	31/12/2021

Executive summary

This document describes the trial tests performed for all the technical components for Risk Assessment (WP4) and Risk Mitigation (WP5) assembled and integrated into the ProTego Toolkit (WP6).

The majority of acceptance and usability tests have been carried out in the testbeds prepared by Ospedale San Raffaele and Marina Salud hospitals, integrating the toolkit into their respective case studies. Some of the tests for the non-functional metrics were performed in the hospitals and other tests were carried out in other lab environments.

The trials described in this document are based on the tests described in deliverable D2.3 and the results provide the final conclusions for the project.

Contributors Table

DOCUMENT SECTION	AUTHOR(S)	REVIEWER(S)
I.	Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)
II.	Michele Cantarutti (OSR), Mirco Pagliarani (OSR), Salvador Garcia (MS), Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)
III.	Michele Cantarutti (OSR), Mirco Pagliarani (OSR), Salvador Garcia (MS), Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)
IV.	Michele Cantarutti (OSR), Mirco Pagliarani (OSR), Salvador Garcia (MS), Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)
V.	Michele Cantarutti (OSR), Mirco Pagliarani (OSR), Salvador Garcia (MS), Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)
VI.	Michele Cantarutti (OSR), Mirco Pagliarani (OSR), Salvador Garcia (MS), Luis Carrascal (Inetum)	Dave Singelee (KUL), Antonio J. Gamito (Inetum)

Table of Contents

I. INTRODUCTION 9

I.1. PURPOSE AND ORGANIZATION OF THE DOCUMENT 9
I.2. SCOPE AND AUDIENCE 9

II. INFRASTRUCTURES FOR TESTING ENVIRONMENTS 10

II.1. OSPEDALE SAN RAFFAELE 10
II.2. MARINA SALUD 12

III. DESCRIPTION OF THE TESTS PERFORMED 14

III.1. ACCEPTANCE TEST 14
III.1.1. Acceptance Test 14
III.1.2. Experimental Procedure 18
III.2. NON-FUNCTIONAL METRICS 20
III.2.1. Non-functional metrics 20
III.2.2. Experimental procedure 25
III.3. USABILITY EVALUATION 25
III.3.1. Usability Inspection and Tests 25
III.3.2. Experimental Procedure 27

IV. RESULTS 30

IV.1. ACCEPTANCE TEST 30
IV.2. NON-FUNCTIONAL METRICS 36
IV.3. USABILITY EVALUATION 40
IV.3.1. Usability Inspection 40
IV.3.2. Usability Tests 41

V. CONCLUSIONS 45

V.1. QUESTIONNAIRE FOR HOSPITALS 45
V.2. PROJECT IMPACT 49

VI. APPENDICES 51

VI.1. ACCEPTANCE TEST EVALUATION 51
VI.2. NON-FUNCTIONAL METRICS EVALUATION 122
VI.3. INTEGRATION TOOLKIT NEW FEATURES 149
VI.3.1. Istio 149
VI.3.2. Storage 151

VII. REFERENCES AND INTERNET LINKS 154

Table of Figures

Figure II.1-1 OSR Testing Environment for FoodCoach 11
Figure II.2-1 MS Testing Environment for Pocket EHR 13
Figure III.1-1 Components, objectives, mission, impacts....18
Figure III.1-2 Experimental Procedure Prioritization 19
Figure III.1-3 Acceptance Test Example 20
Figure III.3-1 ASQ Questionnaire 28
Figure III.3-2 SUS Questionnaire 29
Figure VI.2-1 Normal Operation Risk Model....127
Figure VI.2-2 Increased Vulnerability Risk Model....129
Figure VI.2-3 Dependency of Nutritionist's App on FoodCoach Front End 130
Figure VI.2-4 Dependency of Nutritionist on Nutritionist's App 130
Figure VI.2-5: Root Cause Threat 130
Figure VI.2-6: Control Strategy Options 131
Figure VI.2-7: Channel Suspension 133
Figure VI.2-8: Detail of Loss of Timeliness at PatientApp....133
Figure VI.2-9: Software Patched Control Effect....134
Figure VI.2-10 Network Slicing test lab 143
Figure VI.2-11 Network performance test 1....143
Figure VI.2-12 Network performance test 2....144
Figure VI.2-13 Ping test 145
Figure VI.2-14 Ping test statistics....146
Figure VI.2-15 Network Slicing throughput test 147
Figure VI.2-16 Ping test encrypted....148
Figure VI.3-1 Integration Toolkit Traffic Flow....151

List of Tables

Table III.1-1 Acceptance Tests....14
Table III.2-1 Non-functional metrics....20
Table III.3-1 Use cases for Usability Inspection....26
Table IV.1-1 Acceptance Test....30
Table IV.1-2 Success Rate....34
Table IV.2-1 Non-functional metrics....36
Table IV.3-1 Usability Inspection....40
Table IV.3-2 Usability tests scoring for Network Slicing....41
Table IV.3-3 Usability tests scoring for SSM....42
Table IV.3-4 Usability tests scoring for SIEM....43
Table IV.3-5 Usability tests scoring for Continuous Authentication....44
Table V.1-1 Questionnaire for Hospitals about the results....45
Table VI.2-1 Parquet Modular Encryption benchmark tests....135
Table VI.3-1 Istio integration by toolkit component....150
Table VI.3-2 Storage integration....153

Table of Acronyms and Definitions

Acronym	Definition
ACL	Access Control List
API	Application Programming Interface
APP	Application
ASQ	After Scenario Questionnaire
AT	Acceptance Test
AWS	Amazon Web Services
CA	Continuous Authentication
CLI	Command Line Interface
CPU	Central Processing Unit
D-ABAC	Dynamic Attribute-Based Access Control
DBMS	Data Base Management System
EC	European Commission
EDR	Endpoint Detection and Response
EHR	Electronic Health Records
EKS	Elastic Kubernetes Service
ELB	Elastic Load Balancing
EMR	Electronic Medical Records
FHIR	Fast Healthcare Interoperability Resources
GUI	Graphical User Interface
HL7	Health Level 7
IAM	Identity and Access Management
IoT	Internet of Things
IT	Information Technology
JWT	JSON Web Token
KMS	Key Management Service
LAN	Local Area Network

MAC	Media Access Control
MS	Marina Salud
Mxx	Month xx (of the project)
NFS	Network File System
NVMe	Non-Volatile Memory Express
OPA	Open Policy Agent
OSR	Ospedale San Raffaele
RAM	Random Access Memory
RBAC	Role-Based Access Control
SaaS	Software as a Service
SSD	Solid State Drive
SIEM	Security Information and Event Management
SSM	System Security Modeller
SUS	System Usability Scale
TLS	Transport Layer Security
UC	Use Case
URL	Unified Resource Locator
VPC	Virtual Private Cloud
VPN	Virtual Private Network
WP	Work Package

I. INTRODUCTION

I.1. Purpose and organization of the document

This is the fourth and last deliverable (D7.4) released by ProTego WP7 and provides an overview of the final trials. It consists of the result of activities performed in this work package (Task 7.4).

The input for this deliverable was scheduled to be executed from M31 to M36, and the objective of this task was to setup and run trial scenarios and go further in testing and validating new results from the other WPs. This deliverable analyses the experiments carried out in the last iteration of the project.

The deliverable consists of four major parts: the first describes the environments where the tests have been carried out, the second describes the tests performed, the third describes the results found, and the fourth the conclusions for the tests and the whole project.

The appendices describe additional information related to the tests performed that provide a deeper insight on the trials and is considered to have a special relevance.

Based on the tests described in deliverable D2.3, there are three different objectives for the tests performed:

- Evaluate Acceptance test success rate, that is, the fraction of implemented scenarios over the total amount of specified scenarios.
- Evaluate Non-functional success rate, that is, the fraction of acceptable metrics over the total amount of specified metrics.
- Evaluate Usability, that is, the extent to which it can be used by specified users with effectiveness, efficiency, and satisfaction in a specific context of use.

I.2. Scope and audience

This deliverable is considered as a public document to the ProTego project. It provides a description of the status quo of the performed trials.

This deliverable is primarily intended for EC Officers and project-appointed reviewers. Further, the document is relevant for project members.

II. INFRASTRUCTURES FOR TESTING ENVIRONMENTS

This section describes the testing environments that have been used in each case study. Appendix VI.3. Integration Toolkit New Features, adds more technical details about the different deployment in both environments.

II.1. Ospedale San Raffaele

This case study uses an existing application that has been modified to use the ProTego Toolkit. This application is FoodCoach, a food recommendation system that suggests personalized nutrition plans for end-users. It represents a traditional archetypal distributed application that allows out of boundary data exchange, opening up to the usage of personal mobile devices.

The testers installed the ProTego toolkit as a preliminary activity in order to be able to test it. In preparation for the installation process, the necessary infrastructure was setup. The first step involved the creation of four virtual machines according to the provided specifications. In particular, they were created in the same subnetwork, and each of them was given a static IP address. The four virtual machines were used as part of the Kubernetes cluster where all the ProTego Toolkit components run. The roles of the virtual machines are the following:

- Kubernetes master node, where the Control Plane and etcd services run, and no workloads are executed.
- Rancher node, where the Rancher management system runs, and no workloads are executed.
- Two worker nodes, where the workloads for the ProTego toolkit actually run.

The services of the ProTego Toolkit run in Docker containers. Kubernetes was installed in order to orchestrate these containers, and then Rancher was installed to add a complete UI and workload management layer to Kubernetes. By using the Rancher interface, it was possible to deploy the services of the ProTego Toolkit. Once the ProTego toolkit was setup and running, the medical application could use the exposed services. Figure II.1-1 depicts the configured environment used during the testing activities.

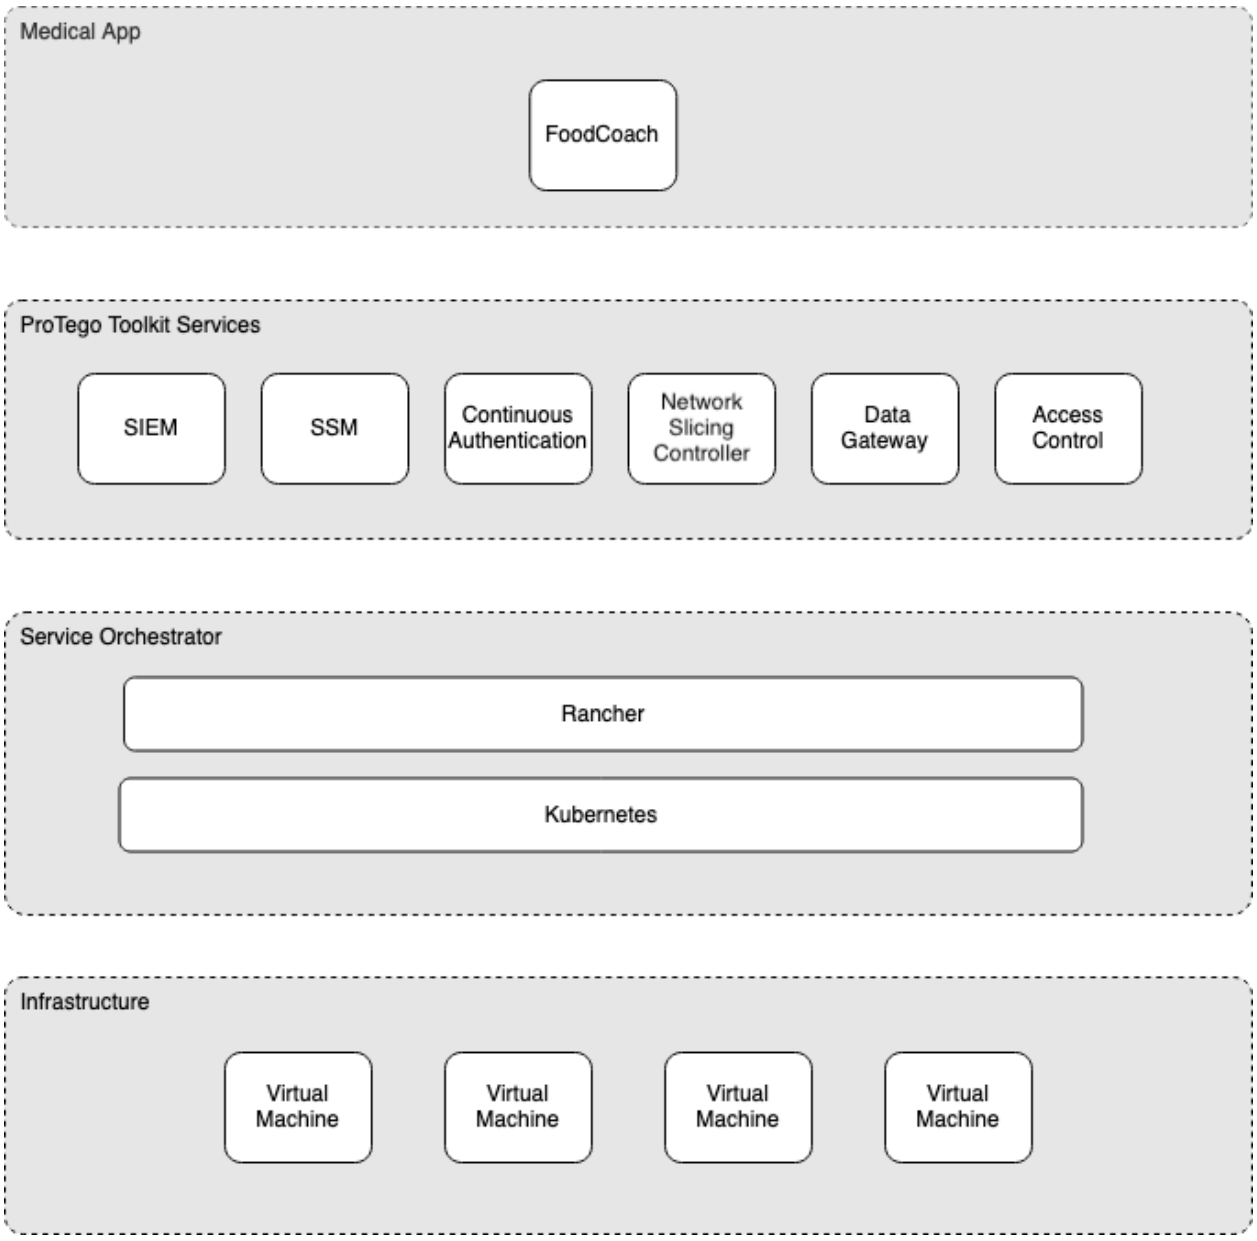

Figure II.1-1 OSR Testing Environment for FoodCoach

II.2. Marina Salud

This case study is more complex than the previous since it involves a hybrid cloud deployment. Therefore, not all the ProTego Toolkit components have been tested in this case. By the work performed, the following components, configuration, and functionalities, have been setup:

1. Medical APP frontend: The Pocket EHR mobile application has been developed, which allows patients to enter health parameters of interest to follow up after a heart failure hospitalization process. The operation of an IoT device has been simulated through the use of a Raspberry Pi device, that sends vital parameters automatically. Pocket EHR, in addition, allows doctors to make complex queries to the Data Gateway, obtaining indicators over all patients on the platform (patients with alert symptoms, etc.)
2. Cloud premises: The interfaces to the Internet (API Gateway) and the backend of the mobile application (Lambda functions) have been developed. A Kubernetes cluster has also been created based on AWS EKS service, and managed by an instance of Rancher. Within the cluster, the pods for all the ProTego toolkit components have been deployed in an isolated way using AWS VPC, so that they are only accessible from Lambda functions, which minimizes the exposure of the ProTego Toolkit components, thus minimizing the attack surface as well. The only exception is Network Slicing that was deployed in the hospital premises.
3. Hospital premises: Integration routes have been developed in Rhapsody, a commercial Integration Engine. They allow to make requests to the ProTego Data Gateway to retrieve data entered by patients using Pocket EHR, and carry out the transformation from FHIR to HL7 2.3, to integrate that data into the patient's EMR in the hospital, implemented through CERNER Millennium Clinical History system, one of the most advanced Medical Record systems in the market. As mentioned before, Network Slicing was deployed in hospital premises.

In this way, it is possible for patients, either through the Pocket EHR mobile APP or through an IoT device, to enter data into the ProTego Data Gateway installed in the cloud environment, and this information can be integrated into the patient's medical record. The hospital premises are not exposed to the Internet, and the ProTego services are isolated and can only be requested from the Lambda functions in the cloud premises.

The image below depicts the use case:

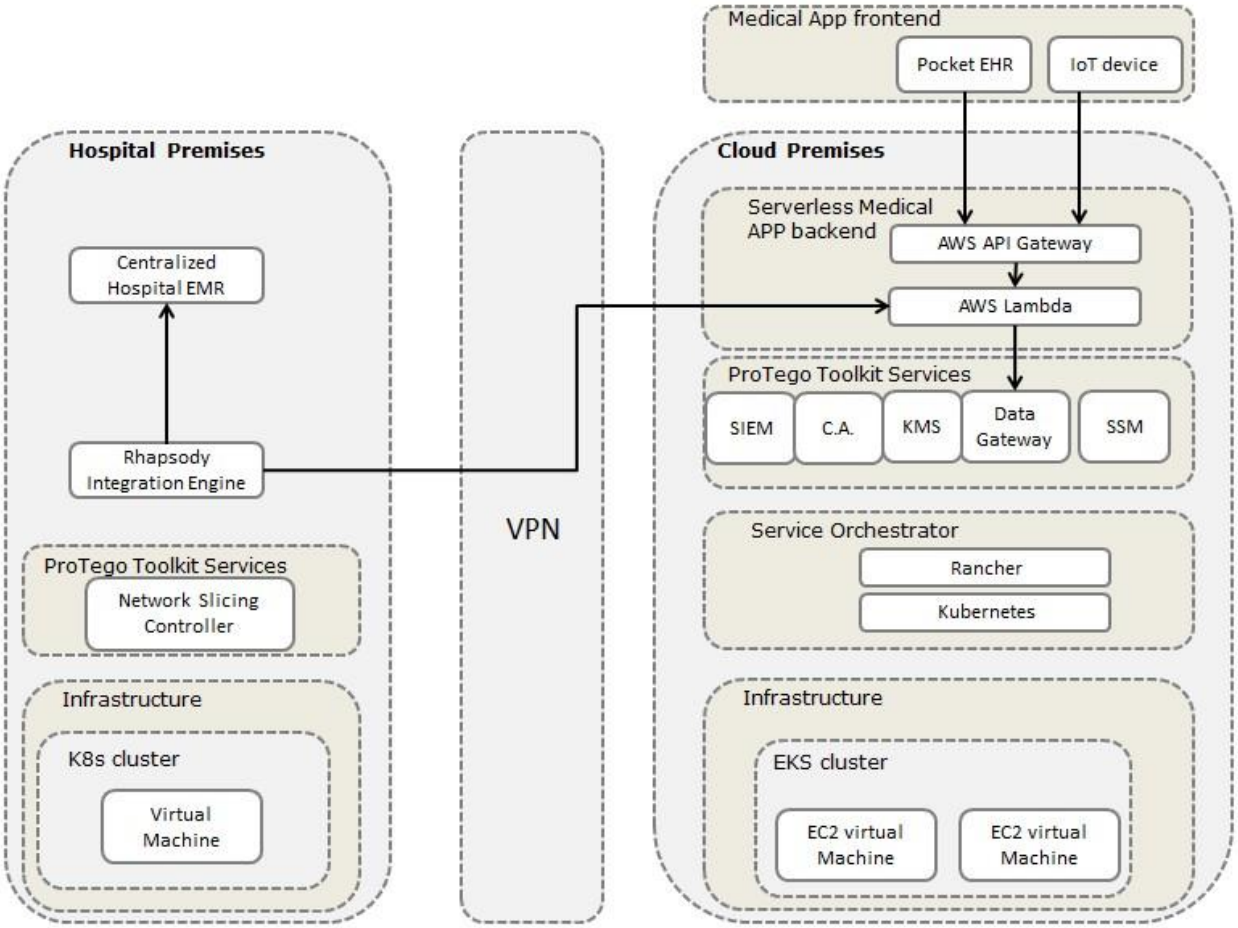

Figure II.2-1 MS Testing Environment for Pocket EHR

Acting as a ProTego adopter, Marina Salud provides the underlying infrastructure and logic elements outside the ProTego Toolkit:

- Cloud provider account, as the repository to all the elements outside the hospital premises.
- Pocket EHR frontend plus emulated IoT device.
- Pocket EHR serverless backend.
- EKS cluster, managed by Rancher.
- Isolation of the EKS cluster (and services) through VPC.
- VPN between the hospital and cloud premises.
- Standard base logic to integrate the ProTego Toolkit managed data into a commercial EMR system.

III. DESCRIPTION OF THE TESTS PERFORMED

This section reports the methods used by testers for the validation of the ProTego toolkit. The first subsection describes the Acceptance tests and the relative experimental procedure carried out in order to execute these tests. Similarly, the second subsection describes the experimental procedure that was employed to collect the Non-functional metrics. The third subsection focuses instead on the description of Usability Inspection, which was adopted to evaluate the system usability.

III.1. Acceptance Test

III.1.1. Acceptance Test

The Acceptance tests constitute the portion of the acceptance criteria utilized to evaluate the implemented functionalities and their associated quality attributes. Each Acceptance test is composed of a set of scenarios. Each scenario describes through an example how a use case should function in the executable software [1]. Deliverable D2.3 section VIII describes these tests in detail.

For convenience, Table III.1-1 reports the Acceptance tests and related scenarios evaluated in this deliverable.

Table III.1-1 Acceptance Tests

Acceptance test	Scenario
AT01: Deploy cluster	S01.1: Operator deploys the cluster
	S01.2: Insufficient authorization
AT02: Install Data Gateway	S02.1: Operator installs Data Gateway
	S02.2: Insufficient authorization
AT03: Install Network Slicing	S03.1: Operator installs Network Slicing
	S03.2: Insufficient authorization
	S03.3: Access point already registered
AT04: Install SIEM	S04.1: Operator installs SIEM
	S04.2: Insufficient authorization
AT05: Install Continuous Authentication	S05.1: Operator installs Continuous Authentication
	S05.2: Insufficient authorization
AT06: Install SSM	S06.1: Operator installs SSM
	S06.2: Insufficient authorization

AT07: Register user	S07.1: Administrator registers a new user
	S07.2: Insufficient authorization
	S07.3: User identifier already taken
AT08: Conduct first-time risk assessment	S08.1: Security operator assesses risks
	S08.2: Insufficient authorization
	S08.3: First-time assessment already conducted
AT09: Assess prospective risks to the infrastructure	S09.1: Security operator assesses application risks
	S09.2: Insufficient authorization
	S09.3: First-time risk assessment is still to be conducted
AT10: Install application	S10.1: System operator installs a new application
	S10.2: Insufficient authorization
	S10.3: Already existing application
AT11: Configure application network slices	S11.1: Security operator configures network slices
	S11.2: Insufficient authorization
	S11.3: Network slices are already configured
AT12: Configure application logging mechanism	S12.1: Security operator specifies logging mechanism
	S12.2: Insufficient authorization
AT13: Specify application access control	S13.1: Security operator specifies access control
	S13.2: Insufficient authorization
AT14: Configure mobile device	S14.1: Security operator configures continuous authentication agent
	S14.2: Mobile device already configured
AT15: Store initial medical data	S15.1: Data operator stores initial medical data
	S15.2: Insufficient authorization
	S15.3: Application initial resources are already stored
AT16: Register mobile device	S16.1: Network operator registers a new mobile device
	S16.2: Insufficient authorization

	S16.3: Mobile device is already registered
AT17: Store medical data	S17.1: Nutritionist stores patients' medical data
	S17.2: Insufficient authorization
	S17.3: Authorization token is expired
	S17.4: Authorization token is ill-formed
AT18: Retrieve medical data	S18.1: Nutritionist queries for patient's data[1]
	S18.2: Insufficient authorization
	S18.3: Medical data have been tampered with
	S18.4: Authorization token is expired
	S18.5: Authorization token is ill-formed
AT19: Assign IoT device to application user	S19.1: Patient is assigned a new IoT device
	S19.2: Assign IoT device to a new user
	S19.3: Insufficient authorization
	S19.4: Authorization token is expired
	S19.5: Authorization token is ill-formed
AT20: Log custom application event	S20.1: Correct application key
	S20.2: Wrong application key
AT21: Send medical data securely	S21.1: IoT device sends some medical data securely
	S21.2: Authorization token is ill-formed
	S21.3: Expired ID_token
	S21.4: Protecting against Replay attacks
	S21.5: Protecting against spoofing
AT22: Report suspicious activity	S22.1: Mobile agent reports suspicious activity
AT23: Respond to alert	S23.1: Security operator responds to an alert
	S23.2: Insufficient authorization
AT24: Review alerts	S24.1: Security operator reviews alerts

[1] In deliverable D2.3 this scenario is called *Nutritionist queries for some patient's weights*

	S24.2: Insufficient authorization
AT25: Review new risk evaluation	S25.1: Security operator reviews changes in the evaluation
	S25.2: Insufficient authorization
	S25.3: First-time risk assessment is still to be conducted
AT26: Reflect infrastructure changes	S26.1: Security operator reflects recent changes
	S26.2: Insufficient authorization
	S26.3: First-time risk assessment is still to be conducted

III.1.2. Experimental Procedure

After the initial setup, the testers followed the experimental procedure reported below:

1. **Prioritization**: In the course of the development of the ProTego Toolkit, some functionalities underwent a shift of priorities, as they became less relevant than initially envisioned. For this reason, the testers assigned to each scenario a priority, in order to highlight those functionalities most relevant to demonstrate the project objectives (shown in Figure III.1-1).

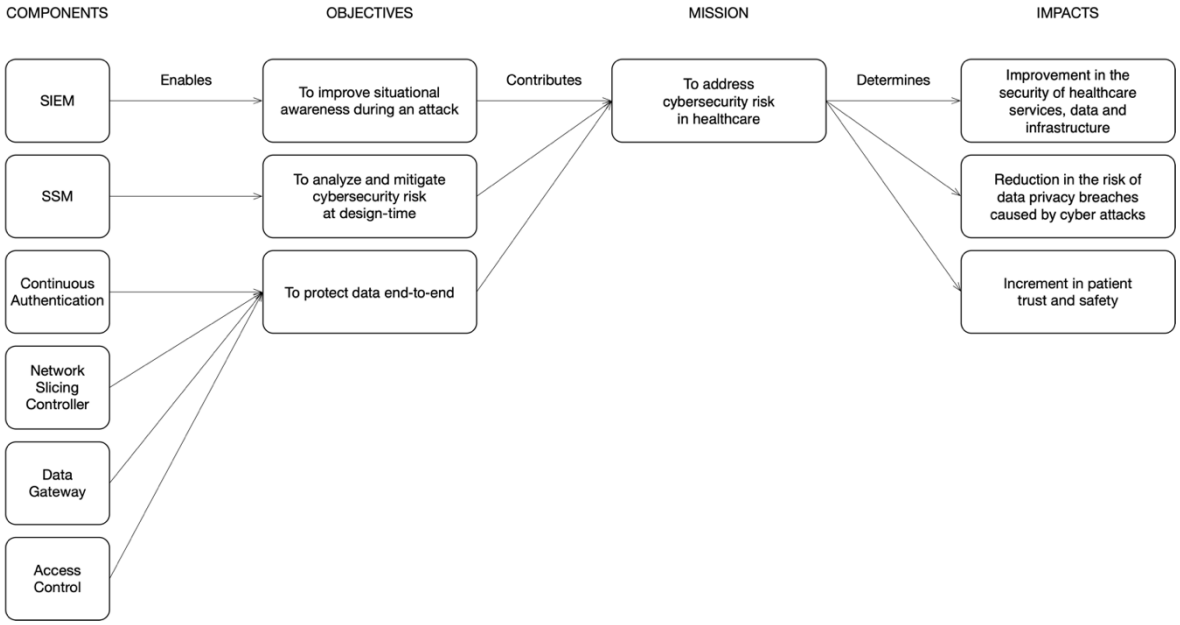

Figure III.1-1 Components, objectives, mission, impacts

Specifically, the testers prioritized the scenarios according to two dimensions: importance and urgency [2]. Every requirement can be considered as being either important to achieving project objectives or not so important, and as being either urgent or not so urgent.

As shown in the Figure III.1-2, these alternatives yield three possible combinations, which have been used to define a priority scale [1]:

- **High priority**: The specified functionalities are both important and urgent to achieve project objectives.
- **Medium priority**: The specified functionalities are important but less urgent to achieve project objectives.
- **Low priority:** The specified functionalities are neither important nor urgent to achieve project objectives. However, these scenarios reflect the functionalities that would be desirable in a market-ready solution.

This priority has helped to establish the order in which the different tests were going to be carried out.

Moreover, a few scenarios have been marked as *Outdated*, and were thus excluded from the evaluation (i.e., AT10 and S19.3 to S19.5). This is because those specific functionalities no longer fit in the ecosystem of the project technologies as originally planned.

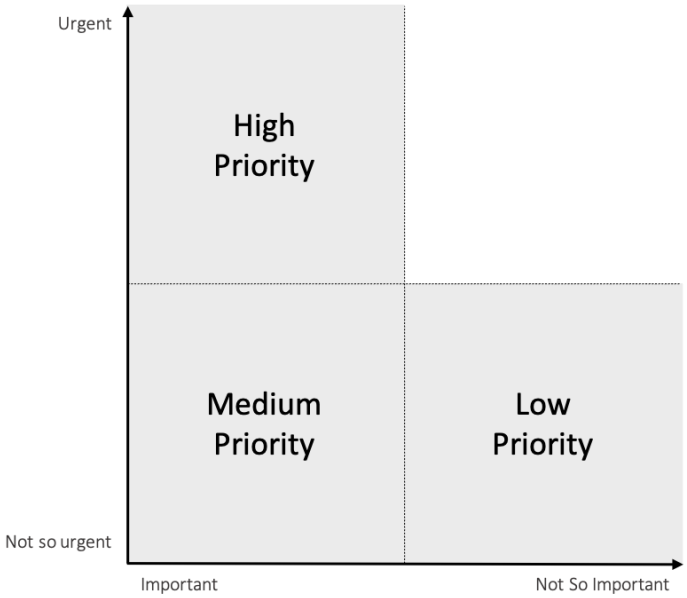

Figure III.1-2 Experimental Procedure Prioritization

2. **Execution**: The testers exercised ProTego functionalities, in order to verify for each scenario, the required satisfaction preconditions, the event or action causing the interaction, and the expected outcome resulting from such an interaction [3].
3. **Evaluation**: The testers marked as *Passed* those scenarios in which the system behaviour corresponded to the behaviour described in the scenario. Otherwise, the testers marked the scenario as *Failed*. Figure III.1-3 shows an example of a scenario that was marked as *Passed* (that is, the successful scenario “Operator deploys the cluster” of the Acceptance Test “Deploy cluster”). Some functionalities of the Toolkit were initially supposed to be provided by a centralized system. Because such a centralized system has not been developed, the responsibility of providing these functionalities was offloaded to the individual components of the Toolkit instead. In the light of this consideration, the following rationale was applied in order to assess the status of the scenarios of these specific cases:
    - If a certain functionality is individually provided by all the components that require it, the corresponding test is marked as *Passed*.
    - If a certain functionality is provided by none of the components that require it, the corresponding test is marked as *Failed*.

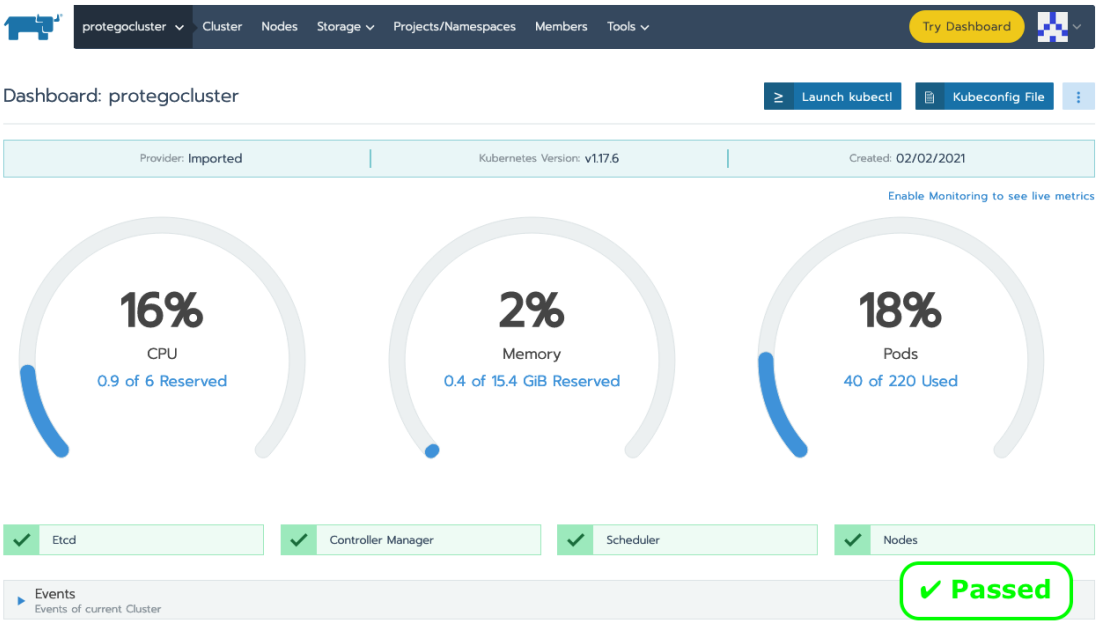

Figure III.1-3 Acceptance Test Example

4. **Analysis**: The testers calculated the success rate, that is, the percentage of *Passed* scenario over the total amount of scenarios. The results have been reported according to the priority scale mentioned above: High priority, Medium priority, and Low priority.

III.2. Non-functional metrics

III.2.1. Non-functional metrics

In order to calculate the Non-functional success rate, a set of metrics that measure non-functional qualities of ProTego were defined, so as to provide an evaluation of specific characteristics of the system, as described in section VIII.2 of Deliverable D2.3. Several questions were posed to determine to what extent the project's objectives were achieved, and each metric helped to answer one of these questions.

For convenience, Table III.2-1 reports the Non-functional metrics evaluated in this deliverable, along with the acceptance criterion of each metric.

Table III.2-1 Non-functional metrics

Questions	Component	Metric	Description	Acceptance criteria
Q1: How effectively is situational awareness improved?	SIEM	Access Control alerts	It describes an alert raised when the Data Gateway's "Authtoken" is not valid	Alert raised from Invalid Authtoken
			It describes an alert raised when the Access Control verifies that signature has expired	Alert raised from Token signature expired
			It describes an alert raised when the	Alert raised from Invalid token

			Access Control verifies that the token is invalid	
			It describes an alert raised when the token type is neither “Keycloak” nor “AWS-cognito”	Alert raised from Wrong token type
			It describes an alert raised when the unwrapped key is not in the base64 encoded format	Alert raised from Wrong data format
			It describes an alert raised when the header structure is not correct	Alert raised from Incorrect header
			It describes an alert raised when the payload structure is not correct	Alert raised from Incorrect payload structure
	SIEM	Continuous Authentication alerts	It describes an alert raised when the trust of the user’s identity drops below a threshold	Alert raised from Untrusted user value
	SIEM	Data Gateway alerts	It describes an alert raised when a cryptographic problem is detected when trying to read a Parquet file	Alert raised from Bad key
	SIEM	Network Slicing alerts	It describes an alert raised when a user tries to perform a malicious scan on an access point	Alert raised from malicious scans
	SIEM	Network Slicing alerts	It describes an alert raised when the traffic coming from a user or an access point is anomalous	Alert raised from anomalous traffic
Q2: How efficiently does the system improve situational awareness during an attack?	SIEM	Event collection rate	It describes the number of collected events within a given timeframe	500 Events/Second per Agent
	SIEM	Event processing rate	It describes the number of processed events within a given timeframe	This value depends a lot on the resources allocated (CPU and RAM), number of

				agents, and events per agent.
	SIEM	Total log storage	It describes the total amount of logs stored	For 100 agents, 80GB/month
	SIEM	Log source count	It describes the number of log sources analyzed	• Application • Data Gateway • Access control • Network slicing • Continuous authentication • SSM • System logs • Network devices
	SIEM	Alerts triggered count	It describes the number of alerts raised within a given timeframe	This value can only be calculated after the SIEM has been working for a time
	SIEM	Resource usage	It describes the resources necessary to utilize the SIEM	CPU: 8 cores RAM: 16 GB Disk: 5 GB
Q3: Is the risk level increased as a result of new detected vulnerabilities?	SIEM / SSM	Increase of risk vector level	Risk level increase in response to a deliberate vulnerability injected into the system	SIEM should detect new vulnerabilities. Overall risk level should increase as a result of the vulnerability being injected.
Q4: Is the risk level reduced as a result of the implemented mitigation strategies?	SSM	Reduction of risk vector level	Risk level reduction in response to mitigating control recommendations to address a detected vulnerability	Overall risk level should reduce from the risk level when a vulnerability is injected as a result of the recommendations made by the SSM.
Q5: How effectively are data-at-rest protected?	Data Gateway	AES key size	It describes the number of bits in a key used by a cryptographic algorithm.	At least 128-bit
Q6: How efficiently does the system perform with the	Data Gateway	Writing overhead	It describes the ratio between performing write operations on encrypted data and unencrypted data	MAX +11% time

data-at-rest protected?	Data Gateway	Reading overhead	It describes the ratio between performing read operations on encrypted data and unencrypted data	MAX +65% time
Q7: How effectively are data-in-use protected?	Access control	False Acceptance Rate (FAR)	It measures the percentage of identification instances in which unauthorized persons are incorrectly accepted.	FAR < 10^(-6)
	Access control	False Rejection Rate (FRR)	It measures the percentage of identification instances in which authorized persons are incorrectly rejected.	FRR < 10^(-3)
	Access control	Access control type	It describes the granularity of access control policies that can be applied.	At least RBAC (Role Based Access Control)
	Access control	Number of user roles supported	It measures the number of user roles that the system support	At least two roles (doctor + patient)
	Access control	Security strength of key storage	It is the number of bits in a key used by a cryptographic algorithm	At least 128-bit security
	Continuous Authentication	False Acceptance Rate (FAR)	It measures the percentage of identification instances in which unauthorized persons are incorrectly accepted.	FAR < 0.5%
	Continuous Authentication	False Rejection Rate (FRR)	It measures the percentage of identification instances in which authorized persons are incorrectly rejected.	FRR < 2%
	Continuous Authentication	Accuracy	It measures the proportion of true positives and negatives to the overall tested data	Accuracy > 98%

	Continuous Authentication	Precision	It expresses how frequently the system correctly produces positive classifications. It is calculated as the ratio of true positive to both true and false positive	Precision > 95%
Q8: How efficiently does the system perform with data-in-use protected?	Access control	Communication overhead per access request	The number of bits (specifically related to the access control solution) that need to be sent to the access control framework by a data consumer in order to access a specific resource	MAX 1 KB
	Access control	Number of IAMs supported	It is the number of identity and access management solutions supported in the access control system	At least 1 IAM supported per hospital within ProTego
Q9: How effectively are data-in-transit isolated?	Network slicing	Number of slices	It is the maximum number of slices supported in the wireless segment. Inside each slice one or more clients can send traffic from one or more applications	Min. 3 slices
Q10: How efficiently does the system perform with the data-in-transit isolated?	Network slicing	Throughput per VPN tunnel	Throughput per VPN tunnel refers to how much data can be transferred from source to destination within a given timeframe for each slice.	Min. 100 Mbps
	Network slicing	Latency	It drives the responsiveness of the network that is how fast each conversation can be had. It is measured as the total round-trip time it takes for a data packet to travel. Latency also drives the maximum throughput of a conversation	Max. 5 ms
	Network slicing	Throughput	It refers to how much data can be transferred from	Min. 10 Mbps

			source to destination within a given timeframe.	
	Network slicing	Packet loss	It is measured as a percentage of packets lost with respect to packets sent.	Max. 0.001%
	Network slicing	VPN additional delay	It is the extra communication delay caused by the encryption/decryption of the data flow in the VPN tunnel	Max. 0.7 ms

III.2.2. Experimental procedure

In the following, the experimental procedure that was used to gather, evaluate, and analyse the Non-functional metrics is described:

1. **Execution:** With a few exceptions, the Non-functional metrics of each component were measured by the partner that was responsible for that specific component, differently from the Acceptance Tests, which were carried out by the two hospitals. The reason for this is that these metrics require very specific knowledge about the component in consideration, and their field of study in general. The results are still well grounded, as most of the metrics evaluate properties that are valid regardless of the underlying infrastructure. For example, in the case of the Continuous Authentication, the metric of the False Acceptance Rate solely depends on the goodness of its algorithm, so such a metric did not necessarily need to be collected in the hospital's infrastructure. Moreover, the University of Alcalá had access to a larger pool of testers (a hundred people) for the experiments, and this would not have been possible or easily manageable for the hospitals. On the other hand, the Throughput metric of the Network Slicing component is a lot more dependent on the infrastructure where it is measured. For this reason, this metric was measured in the hospital's infrastructure, along with the other metrics of the Network Slicing.
2. **Evaluation:** The evaluation of the Non-functional metrics is comparable to that of the Acceptance Tests, as there were only two possible outcomes in each evaluation: either successful or unsuccessful. In the former case the measured metric was marked as *Acceptable*, and in the latter, it was marked as *Insufficient*. This classification was chosen to reflect the more nuanced nature of the Non-functional metrics, since they assess more qualitative aspects. In turn, a metric was considered *Acceptable* only if it was within a pre-defined threshold, known as *Acceptance Criterion*.
3. **Analysis:** The testers calculated the Non-functional success rate, that is, the percentage of *Acceptable* metrics over the total amount of metrics.

III.3. Usability Evaluation

III.3.1. Usability Inspection and Tests

In order to evaluate the usability of ProTego, the testers conducted a usability evaluation by performing a Usability Inspection on some use cases, and a Usability Test [4] on some other use cases. The aim of this evaluation is to measure usability and identify usability problems before making the system available to actual end-users. These evaluations were conducted to validate both the initial setup process and the other use cases specified in D2.3.

The Usability Inspection and Usability Tests followed a similar experimental procedure. The Usability Tests differ from the Usability Inspection as the former involve a greater number of users who perform the tests and they involve an additional questionnaire, known as the System Usability Survey (SUS), which provides a way to evaluate each component of the toolkit as a whole. The choice between these two methods depended on the nature of the use case, and it will be explained in more detail later.

Usability was measured in terms of three metrics: (i) *Effectiveness*, that is, the ability to complete a task; (ii) *Efficiency*, that is, the amount of effort required to complete the task and (iii) *Satisfaction,* that is, the degree to which the user is happy with his or her experience while performing the task.

Table III.3-1 presents an introductory list of the use cases that were tested, along with the method of testing of each use case, that is, either a Usability Inspection or a Usability Test.

Table III.3-1 Use cases for Usability Inspection

ID	Use case	Evaluation method
UC01	Deploy cluster	Usability Inspection
UC02	Install data gateway	Usability Inspection
UC03	Install Network Slicing	Usability Inspection
UC04	Install SIEM	Usability Inspection
UC05	Install SSM	Usability Inspection
UC06	Install Continuous Authentication	Usability Inspection
UC08	Conduct first time risk assessment	Usability Test
UC09	Assess prospective risk to the infrastructure	Usability Test
UC11	Configure application network slices	Usability Test
UC12	Configure application logging mechanism	Usability Test
UC13	Specify application access control	Usability Inspection
UC14	Configure mobile devices	Usability Test
UC16	Register mobile devices	Usability Test
UC23	Respond to alert	Usability Test
UC24	Review alert	Usability Test
UC25	Review new risk evaluation	Usability Test
UC26	Reflect infrastructure changes	Usability Test

This table does not contain all the use cases, but only those whose usability has been analyzed, because a portion of the use cases have been excluded from such testing, as will be explained in more detail later. The table in turn is a reference to Table 107 of deliverable D2.3. The only minor difference is that two further use cases have been excluded from the list, namely “UC10 –

Install application" and "UC15 – Store initial medical data". The former use case has been marked as *Outdated*, whereas the latter has been classed as *Automated*.

III.3.2. Experimental Procedure

1. **Task selection:** In this phase, starting from the list of all the use cases, the testers excluded some of them from the testing because they belonged to either of these two categories:
    a. *Automated*: The actor of the use cases is not a human being, but rather an automated agent. For this reason, it is not possible to assess the task usability according to a user's perspective.
    b. *Outdated*: The task is no longer relevant, as previously discussed in the description of the Acceptance Tests.

2. **Test design:** Due to the restrictions caused by the ongoing pandemic, the test has been designed to be run remotely. This has been done by creating an online version of both surveys and by asking the user to perform the tasks by sharing the screen and by turning on the camera during video calls. Some exceptions have been made depending on the use cases to be performed. In these specific cases, the researchers were in the same room as the testers, so that they could observe them and take notes while the tasks were being performed. As mentioned before, two different methods of testing were designed: Usability Inspection and Usability Tests. This is because to analyse and evaluate the usability through usability tests it is necessary to repeat the same test multiple times with different users, and some of the use cases required a specific environment or setting that was difficult to reproduce or that requires a lot of resources. An example of this is "UC01-Deploy cluster", which requires a lot of resources (e.g., virtual machines) and a long process to install and set up the whole cluster. For this reason, testing this type of use case multiple times with multiple people would be cumbersome. Therefore, for these specific use cases a Usability Inspection was conducted instead, so that fewer people were involved (i.e., only those who had already performed the tasks), and only one questionnaire was submitted (instead of two).

3. **User recruitment/engagement:** The evaluation requires a particular category of users, that is, users with an IT background, preferably in cyber security. In some specific cases, the tasks required skills such as being able to read or edit files from the terminal. Therefore, users have been selected from both hospitals, recruiting employees with such experience and knowledge, for example from the IT department. While some testers already had an introductory knowledge of the toolkit as they had been previously involved during other activities of the project, others were introduced to the toolkit for the first time during the tests.

4. **Execution:** The tester tested all the use cases of each component separately. The tester responded to the After-Scenario Questionnaire [5] after performing the Use Cases' tasks. The questionnaire consists of the three following statements:
    - Q1 "I am satisfied with the ease of completing the tasks in this scenario.''
    - Q2 ''I am satisfied with the amount of time it took to complete the tasks in this scenario.''
    - Q3 ''I am satisfied with the support information (online help, messages, documentation) when completing the tasks.''

    Each statement is accompanied by a 7-point rating scale, from ''strongly disagree'' to ''strongly agree,'' as shown in Figure III.2-1 . The questionnaire allows to determine the usability metrics of each use case, namely: *effectiveness* (Q1), *efficiency* (Q2), and

support (Q3). The last metric, that is *satisfaction*, is determined as the average of the other three metrics.

		1	2	3	4	5	6	7	
1. Overall, I am satisfied with the ease of completing the tasks in this scenario	strongly disagree	○	○	○	○	○	○	○	strongly agree
2. Overall, I am satisfied with the amount of time it took to complete the tasks in this scenario	strongly disagree	○	○	○	○	○	○	○	strongly agree
3. Overall, I am satisfied with the support information (online-line help, messages, documentation) when completing the tasks	strongly disagree	○	○	○	○	○	○	○	strongly agree
		1	2	3	4	5	6	7	

Figure III.3-1 ASQ Questionnaire

In the case of a Usability Inspection, the evaluation ended once the tester responded to the ASQ Questionnaire. In the case of a Usability Test, after responding to the ASQ questionnaire, the user was presented with an additional questionnaire, that is the System Usability Scale (SUS) Questionnaire. This is because the SUS was necessary to collect a more general opinion about a component, after the user completed all the tasks related to that component. The SUS questionnaire is composed of ten statements accompanied by a 5-point rating scale (Figure III.32), from strongly disagree to strongly agree and, similarly to the ASQ, those statements evaluate the metrics previously cited: (I) efficiency, (ii) effectiveness, (iii) support and (iv) satisfaction. In conclusion, the SUS is a useful tool to provide additional and more detailed feedback, not only on a specific use case, but rather on a component as a whole.

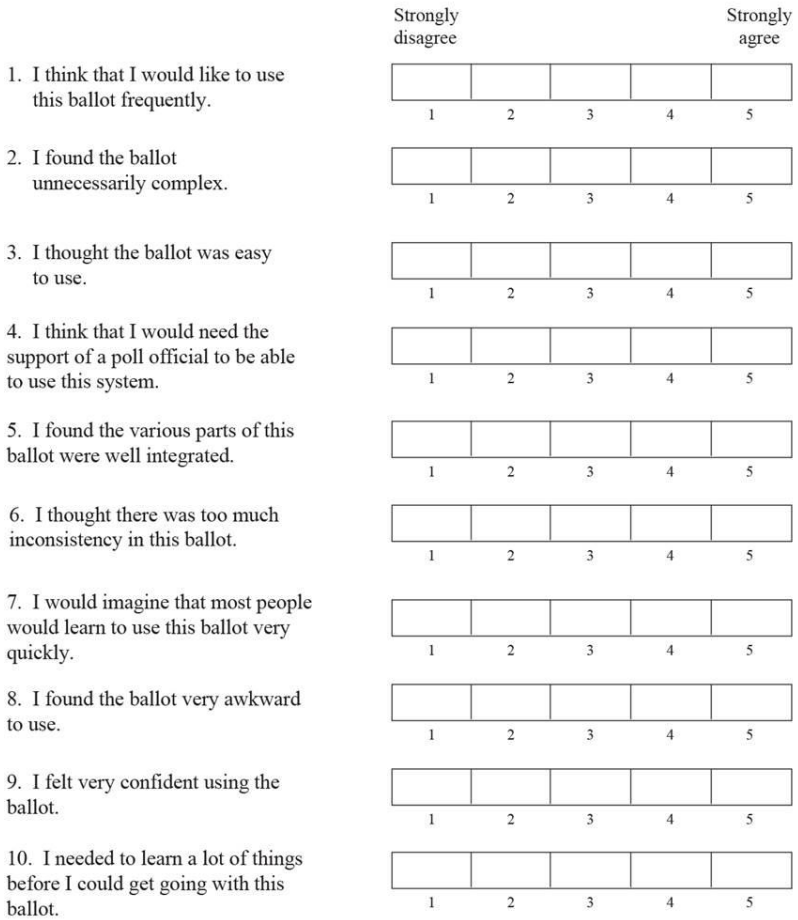

Figure III.3-2 SUS Questionnaire

5. **Evaluation:** In the case of the ASQ questionnaire, the effectiveness, efficiency, and support scores were reported for each examined use case. Satisfaction was calculated as the average score of Q1, Q2 and Q3. In the case of the SUS questionnaire, a usability score was calculated depending on the answers of the tester and resulted in a number between 0 and 100. In general, a score below 50 is considered not acceptable. More specifically, the result of the SUS can be categorized depending on the range the score falls into. The categories are as follows: (I) best imaginable with a score of 100, (ii) excellent with a score between 85 to 99, (iii) good with a score between 68 to 84, (iv) OK with a score between 50 to 68, (v) poor with a score between 30 to 50 and (vi) worst imaginable with a score below 30. [6]

   In addition, the testers provided a set of qualitative suggestions with the aim of providing useful insights to improve usability for the final release.

IV. RESULTS

This section reports a summary of the results concerning the Acceptance Tests, the Non-functional metrics, and the Usability testing. For a more detailed description of the former two, refer to appendices VI.1. Acceptance Test Evaluation and VI.2. Non-functional metrics Evaluation.

IV.1. Acceptance Test

Table IV.1-1 and Table IV.1-2 summarize the results concerning the prioritization and evaluation of the Acceptance Tests and relative scenarios. For a complete description of the results, refer to section VI.1.

Table IV.1-1 Acceptance Test

Acceptance test	Scenario	Priority			Evaluation	
		High	Medium	Low	Passed	Failed
AT01: Deploy cluster	S01.1: Operator deploys the cluster	●			✓	
	S01.2: Insufficient authorization		●		✓	
AT02: Install Data Gateway	S02.1: Operator installs Data Gateway	●			✓	
	S02.2: Insufficient authorization		●		✓	
AT03: Install Network Slicing	S03.1: Operator installs Network Slicing	●			✓	
	S03.2: Insufficient authorization		●		✓	
	S03.3: Access point already registered		●		✓	
AT04: Install SIEM	S04.1: Operator installs SIEM	●			✓	
	S04.2: Insufficient authorization		●		✓	
AT05: Install Continuous Authentication	S05.1: Operator installs Continuous Authentication	●			✓	
	S05.2: Insufficient authorization		●		✓	
AT06: Install SSM	S06.1: Operator installs SSM	●			✓	
	S06.2: Insufficient authorization		●		✓	
AT07: Register user	S07.1: Administrator registers a new user	●			✓	

	S07.2: Insufficient authorization		●		✓	
	S07.3: User identifier already taken		●		✓	
AT08: Conduct first-time risk assessment	S08.1: Security operator assesses risks	●			✓	
	S08.2: Insufficient authorization		●		✓	
	S08.3: First-time assessment already conducted		●			✗
AT09: Assess prospective risks to the infrastructure	S09.1: Security operator assesses application risks	●			✓	
	S09.2 Insufficient authorization		●		✓	
	S09.3: First-time risk assessment is still to be conducted		●		✓	
AT10: Install application	S10.1: System operator installs a new application					
	S10.2: Insufficient authorization					
	S10.3: Already existing application					
AT11: Configure application network slices	S11.1: Security operator configures network slices	●			✓	
	S11.2: Insufficient authorization		●		✓	
	S11.3: Network slices are already configured		●		✓	
AT12: Configure application logging mechanism	S12.1: Security operator specifies logging mechanism	●			✓	
	S12.2: Insufficient authorization		●		✓	
AT13: Specify application access control	S13.1: Security operator specifies access control	●			✓	
	S13.2: Insufficient authorization		●		✓	
AT14: Configure mobile device	S14.1: Security operator configures continuous authentication agent	●			✓	
	S14.2: Mobile device already configured		●			✗
AT15: Store initial medical data	S15.1: Data operator stores initial medical data			●	✓	
	S15.2: Insufficient authorization			●	✓	
	S15.3: Application initial resources are already stored			●	✓	
AT16: Register mobile device	S16.1: Network operator registers a new mobile device	●			✓	

	S16.2: Insufficient authorization		●		✓	
	S16.3: Mobile device is already registered		●		✓	
AT17: Store medical data	S17.1: Nutritionist stores patients' medical data	●			✓	
	S17.2: Insufficient authorization		●		✓	
	S17.3: Authorization token is expired		●		✓	
	S17.4: Authorization token is ill-formed		●		✓	
AT18: Retrieve medical data	S18.1: Nutritionist queries for patients' data[2]	●			✓	
	S18.2: Insufficient authorization		●		✓	
	S18.3: Medical data have been tampered with	●			✓	
	S18.4: Authorization token is expired		●		✓	
	S18.5: Authorization token is ill-formed		●		✓	
AT19: Assign IoT device to application user	S19.1: Patient is assigned a new IoT device	●			✓	
	S19.2: Assign IoT device to a new user		●		✓	
	S19.3: Insufficient authorization					
	S19.4: Authorization token is expired					
	S19.5: Authorization token is ill-formed					
AT20: Log custom application event	S20.1: Correct application key	●			✓	
	S20.2: Wrong application key	●			✓	
AT21: Send medical data securely	S21.1: IoT device sends some medical data securely	●			✓	
	S21.2: Authorization token is ill-formed		●		✓	
	S21.3: Expired ID_token		●		✓	
	S21.4: Protecting against Replay attacks		●		✓	
	S21.5: Protecting against spoofing		●		✓	
AT22: Report suspicious activity	S22.1: Mobile agent reports suspicious activity	●			✓	
	S23.1: Security operator responds to an alert	●			✓	

[2] In deliverable D2.3 this scenario is called *Nutritionist queries for some patient's weights*

AT23: Respond to alert	S23.2: Insufficient authorization		●		✓	
AT24: Review alerts	S24.1: Security operator review alerts	●			✓	
	S24.2: Insufficient authorization		●		✓	
AT25: Review new risk evaluation	S25.1: Security operator reviews changes in the evaluation	●			✓	
	S25.2: Insufficient authorization		●		✓	
	S25.3: First-time risk assessment is still to be conducted		●		✓	
AT26: Reflect infrastructure changes	S26.1: Security operator reflects recent changes	●			✓	
	S26.2: Insufficient authorization		●		✓	
	S26.3: First-time risk assessment is still to be conducted		●		✓	

Table IV.1-2 Success Rate

	# Scenarios	✓ Passed	✕ Failed	✓ Success Rate
Low Priority	3	3	0	100%
Medium Priority	37	35	2	94%
High Priority	26	26	0	100%
Total	66	64	2	97%

The results show that 100% of the high priority Acceptance Tests passed, meaning that all of the most important functionalities expected by ProTego have been successfully implemented. Concerning the medium priority scenarios only two have failed, which is 6%, but they are just secondary scenarios, that may not even be necessary for a final product. This gives us a total of 97% success rate which is satisfactory for the whole project.

The SSM has been evaluated with the following acceptance tests: AT06, AT08, AT09, AT25, AT26. The tests show that the component has the capability of modelling a system (such as the OSR and MS infrastructures) and analysing it to identify potential threats and determine possible mitigation strategies. For example, in the case of OSR, the analysis showed that adding a password strength test to the Frontend service would decrease the potential threats related to this part of the infrastructure. Moreover, the tool allows the user to customise the risk analysis in several ways, for example by assigning a custom level of impact to each adverse effect that results from a threat. Lastly, the SSM is also able to recalculate the risks at run-time, thanks to its integration with the SIEM, which can communicate newly detected vulnerabilities to it via REST APIs that are exposed by the SSM.

The Network Slicing has been evaluated with the following acceptance tests: AT03, AT11, and AT16. The tested scenarios demonstrated that the component could be successfully deployed in the hospitals' infrastructures, and that the component allowed to register devices. Once devices were registered, they could connect to the hospital's network that had been set up for the tests. The tests showed that it was possible to split the network into multiple slices which allowed to isolate network traffic, and which were individually configured to adjust their capacity as needed.

The Access Control & Key management has been evaluated through AT13. The test demonstrated that the component has the capability to specify the application access control by manually intervening on the policy in version 6.1. The access control rules are correctly enforced so that the component can control the access to the medical data. Version 5, which offers a less fine-grained control, was then deployed, and integrated within the healthcare infrastructure. Its main functionality has also been indirectly demonstrated by the fact that the component was successfully used to manage the access to the data stored in the Data Gateway.

The Data Gateway has been evaluated with the following acceptance tests: AT02, AT15, AT17, and AT18. The results show that the Access control and the Data gateway are properly integrated demonstrating their capability of storing and retrieving the patient's records in ProTego. Indeed, the requests were handled correctly, thus accepting the authorized ones and rejecting those unauthorized. This functionality proved useful for the hospitals in preserving health data protection at rest.

The SIEM has been evaluated with the following acceptance tests: AT04, AT12, AT20, AT22, AT23, AT24. The tested scenarios demonstrated that the component is able to monitor application logs, and that it is also able to combine the information gathered from those logs in order to infer a more comprehensive understanding of the nature of the threats to the system. This greatly improves run-time awareness during an attack situation. Moreover, some of the tests demonstrated the integration between the SIEM and the other tools of the toolkit. For example, the SIEM can receive alerts from the Continuous Authentication, and it can detect and send new risks to the SSM.

The Continuous Authentication component has been evaluated through AT05, AT14, and AT22. First, the tests demonstrated the possibility of deploying the Continuous Authentication backend in the hospital testbeds, along with the authentication agent on a mobile device. Then, AT22 demonstrated that the agent can monitor the device to evaluate the level of trustworthiness of a user, and when this level drops below a certain threshold, it can send an alert to the SIEM.

The integration of IoT technology and ProTego data protection tools has been evaluated with the two following acceptance tests: AT19 and AT21. The former focused on the assignment of a device to a user. Contrarily to the original plans, this assignment occurs by providing the credentials of the user, instead of providing a token. For this reason, it no longer makes sense to verify whether the token is valid in the assignment phase, and two scenarios of this test have been marked as *Outdated*. The latter test focused on the use of the device to exchange data with the Data Gateway, and on how data is protected while doing so.

Because the installation of the medical applications occurs independently of the ProTego toolkit, it does not make sense to perform AT10 anymore, and for this reason the test was marked as *Outdated*. On the other hand, the deployment of the cluster and the installation of the whole toolkit was tested with AT01 to AT06.

IV.2. Non-functional metrics

Table IV.2-1 summarizes the results concerning the evaluation of the Non-functional metrics, by reporting whether each metric is *Acceptable* or *Insufficient*, that is, whether its measured value is within its acceptance criterion. For a complete description of the results, refer to section VI.2.

Table IV.2-1 Non-functional metrics

Questions	Component	Metric	Acceptance criteria	Result	Evaluation	
					Acceptable	Insufficient
Q1: How effectively is situational awareness improved?	SIEM	Access control alerts	Alert raised from Invalid AuthToken	Alert raised	✓	
			Alert raised from Token signature expired	Alert raised	✓	
			Alert raised from Invalid token	Alert raised	✓	
			Alert raised from wrong token type	Alert raised	✓	
			Alert raised from Wrong data format	Alert raised	✓	
			Alert raised from Incorrect header	Alert raised	✓	
			Alert raised from Incorrect payload structure	Alert raised	✓	
	SIEM	Continuous authentication alerts	Alert raised from Untrusted user value	Alert raised	✓	
	SIEM	Data gateway alerts	Alert raised from Bad key	Alert raised	✓	
	SIEM	Network Slicing	Alert raised from malicious scans	Alert raised	✓	
	SIEM	Network Slicing	Alert raised from anomalous traffic	Alert raised	✓	
Q2: How efficiently does the system improve situational	SIEM	Event collection rate	500 Events/Second per Agent	The SIEM is able to collect 500 events per second per agent	✓	

awareness during an attack?		Event processing rate	This value depends a lot on the resources allocated (CPU and RAM), number of agents, and events per agent.	A bit slower than expected but no events dropped, nor messages discarded (see details in the Results section)	✓	
		Total log storage	For 100 agents, 80GB/month	161GB/month		✗
		Log source count	• Application • Data Gateway • Access control • Network slicing • Continuous authentication • SSM • System logs Network devices	Events collected from all the expected logs	✓	
		Alerts triggered count	This value can only be calculated after the SIEM has been working for a time	4454 security events collected from 6 agents after 3 days of testing	✓	
		Resource usage	CPU: 8 cores RAM: 16 GB Disk: 5 GB	The SIEM uses more resources than initially foreseen		✗
Q3: Is the risk level increased as a result of new detected vulnerabilities?	SIEM / SSM	Increase of risk vector level	The SIEM should detect new vulnerabilities. Overall risk level should increase as a result of a vulnerability being injected.	Risk level increased	✓	
Q4: Is the risk level reduced as a result of the implemented mitigation strategies?	SSM	Reduction of risk vector level	Overall risk level should reduce from the risk level when a vulnerability is injected as a result of the recommendations made by the SSM.	Risk level reduced	✓	
Q5: How effectively are data-at-rest protected?	Data Gateway	AES key size	At least 128-bit	128 bits	✓	
	Data Gateway	Writing overhead	MAX +11% time	+10.9% time	✓	

Q6: How efficiently does the system perform with the data-at-rest protected?	Data Gateway	Reading overhead	MAX +65% time	+65.8% time	✓	
Q7: How effectively are data-in-use protected?	Access control	False Acceptance Rate (FAR)	FAR < 10^(-6)	1.9*10^-34	✓	
	Access control	False Rejection Rate (FRR)	FRR < 10^(-3)	0	✓	
	Access control	Access control type	At least RBAC (Role Based Access Control)	D-ABAC	✓	
	Access control	Number of user roles supported	At least two roles (Doctor and Patient)	Doctor and Patient	✓	
	Access control	Security strength of key storage	At least 128-bit security	256 bits	✓	
	Continuous Authentication	False Acceptance Rate (FAR)	FAR < 0.5%	0.36%	✓	
	Continuous Authentication	False Rejection Rate (FRR)	FRR < 2%	1.63%	✓	
	Continuous Authentication	Accuracy	Accuracy > 98%	99.06%	✓	
	Continuous Authentication	Precision	Precision > 95%	98.43%	✓	
Q8: How efficiently does the system perform with data- in-use protected?	Access control	Communication overhead per access request	MAX 1 KB	817 B	✓	
	Access control	Number of IAMs supported	At least 1 IAM supported per hospital within ProTego	Two IAMs supported	✓	
Q9: How effectively are data-in-transit isolated?	Network slicing	Number of slices	Min. 3 slices	3	✓	
Q10: How efficiently does the system perform with the data-in-transit isolated?	Network slicing	Throughput per VPN tunnel	Min. 100 Mbps	~300 Mbps	✓	
	Network slicing	Latency	Max. 5 ms	~4.3 ms	✓	
	Network slicing	Throughput	Min. 10 Mbps	~14 Mbps	✓	
	Network slicing	Packet loss	Max. 0.001%	0%	✓	

	Network slicing	VPN additional delay	Max. 0.7 ms	+ 0.2 ms	✓	

Q1 and Q2 allowed to better assess the SIEM. In particular, Q1 focused on the integration aspects, that is, how well-integrated the SIEM is with the rest of the components. The SIEM plays a pivotal role in this since it is the main monitoring tool of the whole toolkit. For this reason, it must be able to receive several types of alerts from all the other components, as demonstrated by Q1. On the other hand, the goal of Q2 was addressing metrics that were more related to the resource-consumption aspects of the SIEM. Almost all the measured metrics were acceptable, except for two of them (i.e., the total log storage and the resource usage), as the value of both these metrics turned out to be greater than initially expected. Further details can be found in section VI.2.

A very detailed analysis of the SSM was performed through Q3 and Q4, which aimed at understanding whether the risk level would change depending on different situations. Both questions were answered in an empirical way by running a trial that represented the everyday use of the SSM. In the first portion of the trial, new vulnerabilities were deliberately injected into the system, which resulted in an increase of the risk level in response to the detection of such vulnerabilities, as expected. In the second portion of the trial, the recommendations made by the SSM to reduce the overall risk level were implemented, and the system reported a decrease of the risk level, as expected.

Q5 and Q6 provided a qualitative analysis of the Data Gateway. Most importantly, the provided analysis was useful to assess the overhead of data encryption during read and write operations. The results reported in the table represent the worst-case scenarios, but they can be greatly improved, as explained in section VI.2. In either case, the assessment demonstrated that the impact of data encryption is not particularly significant.

Q7 addressed the extent to which data is effectively protected while in use. This mainly involved two components, namely: the Access Control and the Continuous Authentication. Both the components have an acceptable rate of false positives and false negatives. Moreover, the former is able to offer a type of access control (i.e., D-ABAC) that is even superior to what was originally envisioned (i.e., RBAC), with an acceptable number of roles. In a more nuanced analysis, accuracy and precision were assessed for the latter component, and the analysis revealed they were also within the acceptable criteria. Similarly, the former component was also assessed by means of Q8, in order to determine its efficiency. The overhead per access request was the most significant metric in this case, and its measured value demonstrated that it was below the expected threshold.

Q9 and Q10 provided an insight into the qualitative aspects of the Network Slicing. These measurements may vary depending on the underlying infrastructure. More specifically, these experiments were run in the hospital's infrastructure, in an environment mainly made up of a virtual machine (4 cores and 4GB of RAM), and a Wireless Termination Point. The experiments demonstrated that the Network Slicing meets all the expectations, as it allows creating at least 3 slices, where each slice is able to maintain a Throughput that is within the proposed acceptance criterion. Making the traffic pass through these slices and encrypting the traffic proved to have a limited impact, which was below the acceptable threshold.

In conclusion, the resulting Non-functional success rate is 95%, which is a satisfactory result, since 37 out of the 39 measured metrics turned out to be acceptable, whereas only two of them were insufficient. The insufficiency of two metrics was due to the initial underestimation of the necessary resources to run the SIEM. This is not an issue per se, but rather it provides a good indication of the resource requirements for any future use of the toolkit.

IV.3. Usability evaluation

The two following sections summarize the results that were obtained during the Usability Inspection and the Usability Tests.

IV.3.1. Usability Inspection

The following table summarizes the evaluation concerning the Usability Inspection. This table lists approximately half of the use cases whose usability was evaluated, as the other half was investigated through the Usability Tests, which are reported in the next section.

Table IV.3-1 Usability Inspection

TYPOLOGY	ID	TASK NAME	EVALUATION			
			Effectiveness	Efficiency	Support	Satisfaction
Installation	UC01	Deploy cluster	5,8	5,2	6,2	5,7
	UC02	Install Data Gateway	5	5,2	4	4,7
	UC03	Install Network Slicing	5,4	5,4	6	5,6
	UC04	Install SIEM	5,4	5,8	5	5,4
	UC05	Install SSM	6	6	5,4	5,8
	UC06	Install Continuous Authentication	6,2	6,2	5,8	6
Configuration	UC13	Specify application access control	6	5,4	5,6	5,6

As can be seen, most of these use cases are related to the installation of the toolkit and of its components, except for UC13, which was simply included because it has no graphical interface, and its usability cannot be assessed in a very detailed manner. The deployment of the cluster itself (UC-01) is an onerous operation, as it requires a lot of resources, and allocating these resources several times for multiple testers would be impractical, and this the main reason why it was decided to perform a usability evaluation only with the people who had been exposed to this process. Moreover, the initial installation of the cluster requires downloading and installing tools such as Docker, Kubernetes, and Rancher, which is inevitably a long process that cannot be avoided. The deployment of the cluster is clearly one of the most important tasks, as it lays the foundations of the whole toolkit, and it would not be possible to deploy any other component without it. Although the process can be lengthy, as demonstrated by the fact that the *efficiency* is the metric with the lowest score in UC-01, the consensus among the hospitals is that excellent support was provided for it, as demonstrated by the high score of the *support* metric.

Similar arguments can be applied to the installation of the components of the toolkit (UC02 to UC06). Installing each component just once is neither particularly time-consuming, nor resource-intensive, though installing and uninstalling each component several times can be somewhat cumbersome, for example because of the creation of static resources such as persistent volumes. Once again, this justifies the decision to perform a usability inspection for these use cases instead. Modularity has always been one of the goals of the toolkit since its inception, and therefore each

component can be installed individually and was tested individually. The Rancher platform provides a unified interface for the installation of all the components, but smaller differences persist, as each component has its own initial configuration that requires slightly different steps. The installation of each component was evaluated individually, and the results show that each partner managed to provide support for it to varying degrees of completeness, while focusing on making the process easier (i.e., more *effective*) or faster (i.e., more *efficient*) depending on their possibilities. By comparing the results of this evaluation with that done at the time of D.7.3, it stands out that there has been an improvement in the obtained results related to the installation of the components. Several factors contributed to this change, and they are mostly related to the changes that have been made to the platform and Rancher during the last stage of the project. These changes include the adoption of Istio, which facilitates the use of services, and Longhorn, which facilitates the use of persistent volumes. The installation process was further simplified by including default configurations for many of the components.

While it would be possible to improve the installation process of the toolkit, the testers reported that receiving direct support from each partner during the installation is what had the most positive and significant impact on the experience. In conclusion, it is recommended that a consultant supervise the user during the installation, as this can simplify things and lead the user to a more immediate use of the toolkit. This is not a negative consideration, but rather a realistic one, as the ProTego toolkit requires a strong background in many highly specialized fields.

IV.3.2. Usability Tests

In the following, the results of the Usability Tests are reported.

As stated before, users with a background in IT were recruited for the tests (e.g., from the IT department of the hospital) and it is important to note that, differently from the Usability Inspection, the use cases tested via the Usability Tests have been tested repeatedly by more users to gather more insights and feedback on the usability of the components.

It is also worth noting that the goal of the Usability Tests is not only evaluating the usability of the components, but also the quality of the educational material that was provided by each partner. This material was provided in the form of slides that presented how to use a specific component, and it was provided to the testers before and during the usability tests, to guide the tester through the necessary steps. In some cases, the partners produced a shorter version of their material, which served as a tutorial for the specific use cases that had to be tested. The educational material was evaluated through a part of the ASQ and SUS questionnaires and through the collection of direct feedback from the users who made use of it.

As each test was run multiple times with different people, each table in the following only reports the average result of each metric measured by the ASQ questionnaire (namely, *effectiveness*, *efficiency*, and *support*) and the average score obtained through the System Usability Scale (SUS) questionnaire.

IV.3.2.a. Usability tests | Network Slicing

Table IV.3-2 Usability tests scoring for Network Slicing

ID	DESCRIPTION	EVALUATION				SUS
		Effectiveness	Efficiency	Support	Satisfaction	
UC11	Configure application network slices	5,2	5,8	6	5,7	64,5

UC16	Register mobile devices	6,4	6,4	6	6,3	

The results of the tests (Table IV.3-2) show that the testers had a favourable opinion on the usability of the Network Slicing component. Out of the two tested use cases, UC16 was clearly the simplest one, and this is demonstrated by the generally more positive results it received with regard to its *effectiveness* and *efficiency*. To perform the tasks of use case UC-16, the testers could avail themselves of the Network Slicing controller, which offered a graphical interface (5G-EmPOWER), accessible through the browser. This interface received positive feedback from the testers, as it was deemed simple and intuitive to use. The only minor criticism was a flag-shaped icon that had to be clicked in order to access a specific project, and it was suggested to replace the icon with either an arrow or no icon at all. To register a mobile device, it was necessary to use the MAC address of such a device, and the testers noted that retrieving this piece of information may not be too straightforward for the general user. Also, because mobile devices on the market nowadays are very diverse, retrieving the MAC address may entail different steps on different devices. For this reason, it was suggested by the testers to expand the educational material by adding at least a reference to a tutorial on how to do this on the most common devices nowadays.

There are some similarities with UC-11, as the initial steps mainly involved navigating the 5G-EmPower interface, which once again proved to be intuitive to use. Though, the testers noted that this use case required a deeper knowledge of the concepts of this component, which in turn requires investing more time in reading and understanding the educational material. The most important concepts to be understood are those of *quantum* and *airtime*, which both relate to how much priority (i.e., rate of traffic) a specific slice is assigned. Another thing that proved difficult to be understood at the beginning, is the fact that the quantum of a slice is not defined in absolute terms, but rather in relative terms, that is, in relation to the quantum of the other slices. Though, all the testers understood this concept after playing around with the tool for some time. After configuring the network slices through the web interface, further configurations had to be applied to the Privacy Gateway (PG) and to Wireless Termination Point (WTP), to which the users had to connect via SSH, using a terminal. The same exact configuration had to be applied both to the PG and the WTP, and this step felt redundant according to the testers. The testers suggested finding a way to perform this step only once, as duplicating the configuration can increase the chance of a mistake. Moreover, the testers expressed a preference for the use of a graphical interface like in the first portion of the use case.

The final score of the SUS questionnaire is in line with the above-mentioned comments that were made by the testers, as it proves that the usability of the Network Slicing is overall acceptable, while the suggested minor improvement could certainly bring the score even higher.

IV.3.2.b. Usability tests | SSM

Table IV.3-3 Usability tests scoring for SSM

ID	DESCRIPTION	EVALUATION				SUS
		Effectiveness	Efficiency	Support	Satisfaction	
UC08	Conduct first time risk assessment	4,5	5,2	6,5	5,4	57,5
UC09	Assess prospective risks to the infrastructure	4,7	5,2	6,2	5,3	

UC25	Review new risk evaluation	5,5	5,5	6	5,6	
UC26	Reflect infrastructure changes	5,5	5,2	6	5,5	

As can be seen in Table IV.3-3, the SSM received a SUS score that is lower than that of the other components. This does not mean that the SSM itself is impractical to use, but rather the consensus among the users was that the SSM required the biggest investment of time out of all the components in the toolkit. The length of its educational material is clear evidence of this, as the sole manual is around 30 pages long. Getting accustomed to the terminology and the all the interfaces is not an insurmountable challenge, but the testers referred that a potential customer should be aware of the necessary commitment before adopting this tool. Nevertheless, the testers praised the great effort that has been made by IT Innovation to make such a complex tool as usable and accessible as possible, by providing detailed instructions and constant support through several meetings. This is demonstrated by the fact that the *support* metric scored higher than the other metrics in all the use cases. One important takeaway is that it is essential to provide the support of a consultant during the early adoption of the tool, as mentioned before in this deliverable.
Delving deeper into the specific use cases that were tested, it stands out that the first two use cases scored lower than the other two. That is partially because the latter involved the repetition of some steps that were already performed in the former, and the users felt more confident when repeating the same operations a second time. The first step of UC-08 was considered one of the most time-consuming by the testers, as it required to build a system model. Doing this requires a good knowledge of the SSM and of the infrastructure to be modelled, as the user should be able to identify the assets and the relationships within the infrastructure. Although this is an onerous step, the testers agreed that the resulting model was very useful, as it allowed to have a comprehensive view of the infrastructure, and it simplified the tracking of changes to the infrastructure in the future. After building the model and customising the settings, the risk calculation proved to be very straightforward, as it only required clicking a button. The testers reported that learning how to read the vast amount of information on the interface took some practice, but its completeness was greatly appreciated by all the testers. Another aspect that was appreciated was the immediateness with which control strategies could be viewed, as an informative panel opens whenever an asset is selected.

In conclusion, the tests revealed that the SSM is a complex, yet powerful tool, which requires some training but that, once mastered, can provide vast amounts of information.

IV.3.2.c. Usability tests | SIEM

Table IV.3-4 Usability tests scoring for SIEM

ID	DESCRIPTION	EVALUATION				SUS
		Effectiveness	Efficiency	Support	Satisfaction	
UC12	Configure application logging mechanism	6	5	6	5,6	62
UC23	Respond to alert	5,5	5,2	6	5,5	
UC24	Review alerts	6,5	5,5	6	6	

The testers appreciated the fact that the SIEM graphical interface provided a way to perform all the tasks, and this is reflected in the generally positive results reported in Table IV.3-4. Because the graphical interface provides a lot of functionalities, the testers reported that learning how to use it confidently required some training and time, but the interface was nonetheless deemed less complex than that of the SSM. Indeed, the testers praised the simplicity of performing the three uses cases, as none of them required very advanced knowledge in the field. The creation of a SIEM agent was the only step that required some deeper technical skills, but it was an operation that was performed externally, before the first use case was tested. After that, the configuration of the logging mechanism required the creation of a decoder and some custom rules, which could be done by using the XML language. This did not constitute a challenge for any of the testers, as minimal knowledge of the XML language is common among people with an IT background. The combination of decoders and rules allowed to easily create a tool the detection of potentially dangerous situations, with the only requirement that the user understand how to correlate these events.

The two following use cases mostly dealt with reviewing alerts and information that was collected by the SIEM. In particular, being able to receive an alert notification via Microsoft Teams was seen as a great plus of the tool, as it would not be practical to require the user to constantly monitor the SIEM interface. Once a notification was received, the user needed to navigate to a dashboard that displayed all the information regarding that specific alert. Memorizing the exact steps to reach each page in the interface was not immediate for a beginner, but it was not overwhelmingly difficult either. The interface also offered a list of commands that could be automated so that they would be executed as soon as an alert was fired, though understanding the effects of these commands certainly requires a type of knowledge that goes beyond basic IT skills.

Overall, the SIEM received generally favourable reviews from the testers. Like the other tools, it requires some training, but that is almost only limited to the initial configuration, whereas a simple manual proved to be enough for the monitoring-oriented tasks that make up most of the daily use of this tool.

IV.3.2.d. Usability tests | Continuous authentication

Table IV.3-5 Usability tests scoring for Continuous Authentication

ID	DESCRIPTION	EVALUATION				SUS
		Effectiveness	Efficiency	Support	Satisfaction	
UC14	Configure mobile device	6	6,5	6	6,1	70

The results of the surveys clearly show that the Continuous Authentication made a positive impression on the testers. It must be noted that, compared to use cases of the other components, UC-14 was a simple one. It involved installing an Android application on a mobile phone, configuring the keyboard so that it could be monitored, and then granting accessibility permissions so that the mobile could be locked in case it was stolen. These configuration steps could be performed by using the settings that can be found on any Android device, and they were precisely documented in the provided material. Only one suggestion was provided by the testers, that is, providing an installation wizard to make the process even simpler, so that the user will not have to navigate to the settings and so that the process can be more uniform on all the devices. As expected, this component scored high in the SUS questionnaire.

V. CONCLUSIONS

The positive results reported in this deliverable demonstrate that the objectives for the project have been reached. To provide a solid foundation for that statement two additional resources will be used:

- Hospitals' opinion with a Questionnaire.
- Analysing the Project Impact.

V.1. Questionnaire for Hospitals

To draw conclusions about the results of the project and assess impact, deliverable D3.3 section *VI.7 Opinion about the ProTego toolkit*, included a questionnaire that was designed to gather the opinion of the IT staff in the hospitals that participated in all the testing performed.

This is very important because even the best tool or product will not have any impact if it is not adopted by the targeted organizations. The responses to the questionnaire show the anticipated real perception of adoption by those that would be responsible of the resulting product of the project in the hospitals. The following table includes the questionnaire and the answers received from both hospitals in a unified way:

Table V.1-1 Questionnaire for Hospitals about the results

ID	Question / Answer
#1	**Question:** Does the scope covered by ProTego match areas in which there is a lack of cybersecurity controls within your organization? Explain which ones.
	Answer: Yes, ProTego covers cybersecurity aspects that we do not have covered. Some of them were on our roadmap, such as the incorporation of a SIEM tool. But there are others that provide very new possibilities that were not even contemplated: • the possibility of evaluating the system at design time (SSM). • updating this evaluation in real time by integrating SSM and SIEM. • the encryption of data at rest (Data Gateway) and the granularity in the definition of access rights, at the platform level (Access Control). • the protection of user devices preventing impersonation risks in the access to the health data (Continuous Authentication).
#2	**Question:** Do you think ProTego would help to reduce cyber-risks in your organization? Explain why.
	Answer: We believe that ProTego can effectively help reduce cyber-risks. By covering the previously described areas that are not currently being treated in our organizations, attack vectors can be managed in currently unprotected areas. In addition, the educational material addressed to health staff can reduce the risks derived from incorrect human behaviour related to cyber security.
#3	**Question:** Do you think ProTego would increase situational awareness under an attack and improve the response time?

<table>
<tr><td></td><td colspan="4">Yes, mainly through the real-time information offered by the SIEM, we believe that the response to an eventual attack would be drastically reduced, which would undoubtedly help to reduce the impact.</td></tr>
<tr><td>#4</td><td colspan="4">Can you identify features offered by ProTego that apply to each item on the following table? Short description.</td></tr>
<tr><td></td><td colspan="2" rowspan="2">ProTego Toolkit</td><td colspan="3">Data stages</td></tr>
<tr><td></td><td>Data at rest</td><td>Data in transit</td><td>Data in use</td></tr>
<tr><td></td><td rowspan="3">Cyber-Security Dimensions</td><td>Confidentiality</td><td>Fine-grained definition of access grants to the data. Parquet encrypted files prevent data at rest to be read.</td><td>Isolation provided by Network slicing ensures confidentiality of data in transit.</td><td>Hardware Secure enclaves.</td></tr>
<tr><td></td><td>Integrity</td><td>Parquet encrypted files prevent that data at rest to be modified.</td><td>Isolation provided by Network slicing ensures integrity of data in transit.</td><td>Hardware Secure enclaves.</td></tr>
<tr><td></td><td>Availability</td><td>The integration technology applied (k8s, DevOps) allows the use of persistence layers that ensures availability, like NFS or EFS for cloud infrastructure.</td><td>Network slicing ensures the network resources needed to each service.</td><td>Hardware Secure enclaves.</td></tr>
<tr><td>#5</td><td colspan="4">Can you identify any feature offered by ProTego that is not covered by applications already in the market, as far as you know? List them.</td></tr>
<tr><td></td><td colspan="4">Mainly the features offered by the SSM, as it allows to assess the risk level of a system in design time, and also to assess the impact of future changes to the infrastructure.

Also, the real time integration of the SIEM and the SSM, which allows the SIEM to notify not only risks regarding events that have already occurred (by analysing vulnerabilities in infrastructures and applications) but notifying prospective situations if some change introduced in the design of the system that introduce new risks.

Finally, the Continuous Authentication provides the means to identify impersonation risks in the final user devices. There are some tools that allow managing the corporate devices, performs health assessments (assess the installed apps, patching level...) but as far as we know there is not any tool that assess the trustworthiness of the user that is using the device.</td></tr>
<tr><td>#6</td><td colspan="4">Do you think the ProTego toolkit would be suitable to be adopted in your organization? Explain why.</td></tr>
<tr><td></td><td colspan="4">In the previous answers it has been explained what features offered by ProTego are of interest, some of them are even novel and not offered by other tools already in the market. But assessing adoption is a vital point as a healthcare organization is an ecosystem from the IT perspective and new applications should integrate with the pre-existing ones. In that regard we can see interesting features in the ProTego toolkit:

• The integration framework is based on Kubernetes for deployment and orchestration of services. It allows almost any type of underlying infrastructure. It also allows easy scaling-up based on actual use.

• The ProTego toolkit can be deployed in both cloud and or hospital premises, as it has been demonstrated in MS and OSR pilots respectively.

• It bases the authentication in the JWT standard, and delegates authentication services in external IAMs, what allows to integrate the ACL (user management) with the corporate pre-existing one, something mandatory for the majority of organizations.</td></tr>
</table>

<table>
<tr><td></td><td colspan="2">

It uses FHIR as standard for health data format. It allows easy integration with external applications as the format of the data is well-known. It has allowed to integrate the toolkit with the hospital EMR, what is also mandatory, because it minimizes the isolated data silos and enriched the patient's data record, improving the business intelligence performed on that centralized repository. To be more precise, the toolkit has been integrated with Cerner Millennium, one of the top commercial EMRs.
The ProTego toolkit allows the adoption as a modular system, in a way that the organization can decide which components of the toolkit it decides to adopt at each point in time, allowing to delay those that are not of interest or need a deeper assessment or provide certain requirements before its adoption.

All in all, we think that the ProTego toolkit has addressed the adoption requirements in a successful way.
</td></tr>
<tr><td>#7</td><td colspan="2">Which disadvantage[3] can you see in the ProTego toolkit? Explain.</td></tr>
<tr><td></td><td colspan="2">
They are not disadvantages per se, but we can foresee some aspects that may make it difficult for some organizations to adopt a toolkit like ProTego. They are mainly related with the fact that the toolkit is based on some of the newest standards, and most healthcare organizations are not still ready to use them. To be more precise:

the FHIR standard for data exchange is not yet widely adopted. Nowadays HL7 v2 and v3 are the most used and the majority of tools in the market are compatible with these versions. It's not a blocking problem, since custom developments can be easily developed to translate FHIR format from/to other data formats.
most healthcare organizations are not using DevOps technology, so the use of Kubernetes, Istio, etc. may not be affordable for them right now, not because of the cost of the technology in itself, since it is Open-Source, but for the cost in time, effort and other needed resources to tackle the projects involved. Nevertheless, it is undoubtedly the direction the industry is moving towards, and in a few years any organization would be required to use this kind of technology to use updated applications.
Along the same line, the use of JWT for authentication may not be yet implemented in many healthcare organizations.

In summary, the aspects that may prevent healthcare organizations to adopt the ProTego toolkit are related to the fact that it is based in forthcoming standards, and healthcare organizations need to improve their underlying technology and technical skills to face that challenge successfully. The features offered by these novel technologies could be a good incentive to move forward and accept the challenge.
</td></tr>
<tr><td>#8</td><td colspan="2">Please summarize the opinion about the ProTego toolkit from the scope of each of the following roles you are in charge of:</td></tr>
<tr><td></td><td>Network operator</td><td>Network operators manage the network capabilities of the hospital. They need to guarantee that the network meets the necessary quality of service attributes.</td></tr>
<tr><td></td><td></td><td>Through the Network Slicing tool, the ProTego toolkit provides the means for creating secured slices ensuring the confidentiality of the transmitted data. But it also allows a useful capability: assign the desired resources (bandwidth, especially over wireless networks) to each slice, which in turn increases availability. It has some direct applications for hospitals, for example a slice could be created to transmit DICOM data from medical devices to the unified repository, since this is a heavyweight data format and with this feature we could dedicate a portion of the resources, not affecting other services working under the same physical network.</td></tr>
</table>

[3] The original questionnaire in deliverable D3.3 used the word “contra” which is incorrect, that is why it has been changed to “disadvantage” here.

	Data operator	Data operators are in charge of data flows and how they are stored within the hospital infrastructure. Data operators must guarantee that sensitive data is stored with an appropriate level of protection.
		The underlying technology in ProTego increases the security of data at rest providing an additional encryption feature, that is independent of what is provided by the underlying platform.
	Security operator	Security operators take care of the security concerns of the hospital infrastructure. These include controlling access to the hospital services, assessing the risk associated with the infrastructure, and monitoring the data exchange for possible attacks and foreseeable risks.
		The main features provided by the ProTego toolkit from a holistic perspective, that is, how the toolkit could be integrated with pre-existing IT systems, are: • it integrates with the organization's IAM, thus using a unique user database. This minimizes administration tasks avoiding duplicated information in different user repositories, and the risk of having them out of sync. • it offers integration with external applications, including medical devices, which is aligned with the strategy of reducing isolated data silos that are a main handicap of present medical device systems. • it introduces features to reduce risks of impersonation on the end user device. • it allows performing prospective assessments of the system in design time, allowing to assess the impact a potential change could have. • real-time integration of the SIEM with the SSM, which is a novel and powerful tool to minimize risks. • it is aligned with the NIST CSF, a worldwide recognized standard for cybersecurity assessment.
	System operator	System operators are responsible for installing, configuring and managing computer systems in the hospital infrastructure.
		The fact that the ProTego integration toolkit is based on Kubernetes allows for deployment on different platforms, including cloud premises. This versatility makes it easy to adopt from a system operator perspective, as it makes the management of the platform easier and allows scaling the platform requirements accordingly based on the use and demand.
	Administrator	Special role capable of making unrestricted, system-wide changes (e.g., registering accounts for other IT operators).
		The fact that the toolkit delegates the access permissions to the corporate IAM makes it easier to control the access to the data, as previously mentioned. Regarding administrator access to the platform, it is controlled by the base platform in use. For example, in the case of cloud infrastructures, the user definition and permissions defined in the corporate cloud account make it easier to manage the fine-grained access to the components offered by the cloud, allowing to define specific grants for specific maintenance and management actions.

To summarize the main aspects, the responders could identify features that are understood as useful to prevent or mitigate cyber risks in areas that have not been already covered in both hospitals, and even some of them are not provided by other products already in the market.

From a technical perspective, the versatility of the ProTego toolkit has been highlighted as a key factor to facilitate adoption, allowing to deploy it over different base infrastructures.

And from an organizational perspective, the modularity has been identified as useful, since the toolkit allows partial or progressive adoption, allowing each organization to find the right moment to introduce each component, as the organizational or corporate requirements could be met.

V.2. Project Impact

The DoA envisioned certain KPIs to measure project impact. But after being analysed by the project team, they could not be used realistically, in the case of some of them for the intrinsic nature of the defined KPIs, and in the case of others because the initial situation in the hospitals did not allow to have comparison points. So, an explanation of how those KPIs are fulfilled is given next.

Based on the results of the tests reported in this deliverable, the reports described in previous deliverables such as D3.3 D4.2 or D4.3, and the questionnaire for the hospitals in the previous section, the following key impacts for the project can be outlined:

- **Improved security of Health and Care services, data and infrastructures**
    - **Number of training protocols**: ProTego includes an Educational Framework that provides, among other things 4 videos for non-technical users. It also provides guidelines for the procurement of IoMT (Internet of Medical Things) devices. Evidence of this can be found in deliverable D3.3.
    - **Percentage of false positives reduced**: Machine learning techniques have been used to reduce false positives. Complete and flexible definition of rules is also implemented. Evidence of this can be found in deliverables D4.2 and D4.3.
    - **Time needed to find threats and assess risks**: The use of the ProTego toolkit involves time saving in cost and effort, that is the difference between doing things manually or by a human vs being done by a machine automatically. Before ProTego risk assessment is a complete manual process, but ProTego provides the automated process that performs vulnerability scanning and triggers dynamic risk recalculation, which is more accurate and involves no human effort, except reviewing the results and implementing the recommendations, that has to be a human decision. There is no comparison point since there is no previous solution. Having said that, an initial investment for the creation of risk models is still necessary, as well as staff training to make changes in the risk models. Evidence of this can be found in deliverables D4.2 and D4.3.
    - **Number of attacks exposed**: The ProTego toolkit allows more threats to be detected, by configuring many more points of detection. This shows lots of monitoring on different sources that illustrates the potential of detection without the cost of human effort. Section VI.2. Non-functional metrics Evaluation, and specifically Q1, give evidence of many attacks that could not be detected before.
    - **Number of identified system risks addressed**: The great potential to identify risks, and the capability to provide recommendations gives organizations proper decision-support based on vulnerabilities detected. The lifecycle described in deliverable D4.3 illustrates this clearly.
- **Reduced risk of data privacy breaches caused by cyberattacks**
    - **Number of online accesses to the hospital systems from patients/insurance (payers)**: The use of ProTego increases confidence to provide online services. Managers in Healthcare industry in general are hesitant about providing online services. Using solutions like the ProTego toolkit, many organizations may accept the challenge, since the features offered by ProTego minimizes the intrinsic risks on this kind of online services. In the specific case of one of the pilots, Marina

Salud, some online services are started to be offered, taking advantage of the participation in the ProTego project. Based on their estimation, a number of projected patients can be given. The applications in development assisted by ProTego shows the following figures:

- The initial cardiology pilot is made of some 1,000 patients.
- Moving on to chronic patients increases to some 10,000.
- Finally, as the functionality opens up to all patients the number increases to some 100,000.

- **Number of new services available online**: The project has provided lots of learning. The knowledge gained in the project has assisted both hospitals to design two new online services for the public, that is, FoodCoach and PocketEHR.
- **Total number of data privacy mechanisms adopted in the project by use cases (both hospitals)**: During the course of the project both hospitals have adopted all the technical components developed in WP5 in non-operational or non-production systems to support applications. Whether those systems are finally deployed in production is beyond the control or scope of the project. The Educational material developed in WP3 has been used in both hospitals in the digital format, and printed posters have also been used.

- **Increased patient trust and safety**
  - **Level of trust in remote health devices**: Besides the questionnaire for hospitals detailed in the previous section, two crowdsourced surveys were undertaken by IT Innovation, each with approximately 400 respondents using representative sampling. The surveys investigated their attitude to cybersecurity threats, the controls that ProTego provides to address the threats, and which actors are responsible to enacting the controls. The analysis of the results will be incorporated into an academic paper. The following summarizes the preliminary results.

    The most severe threats in the public's opinions were:

    1. My medical records falling into the wrong hands If the hospital gets hacked.
    2. If I lose my phone or laptop, or someone gets my password, someone may get at my medical records.

    The most effective controls in the public's opinion were:

    1. Adding protection (encryption) to all medical data wherever it's stored, sent or viewed, so it can't be tampered with.
    2. Making sure that my medical data are protected and can only be looked at by medical staff.
    3. Having an automatic lock on a phone or computer app, so my data stays safe even if the phone or computer is stolen.
    4. Making sure that medical staff only see the bits of my data that are essential for my treatment rather than all of it.

    Considering the tools ProTego includes:

    1. Control 1 is general encryption and Control 4 is specifically about Parquet encryption.
    2. Control 2 references the Data Gateway.
    3. Control 3 relates to Continuous Authentication.

VI. APPENDICES

VI.1. Acceptance Test Evaluation

AT01: Deploy Cluster

ID: S01.1

Name: Operator deploys the cluster

Priority: High

Evaluation: ✓ Passed

Description: The testers successfully deployed the cluster. The resulting resources can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.

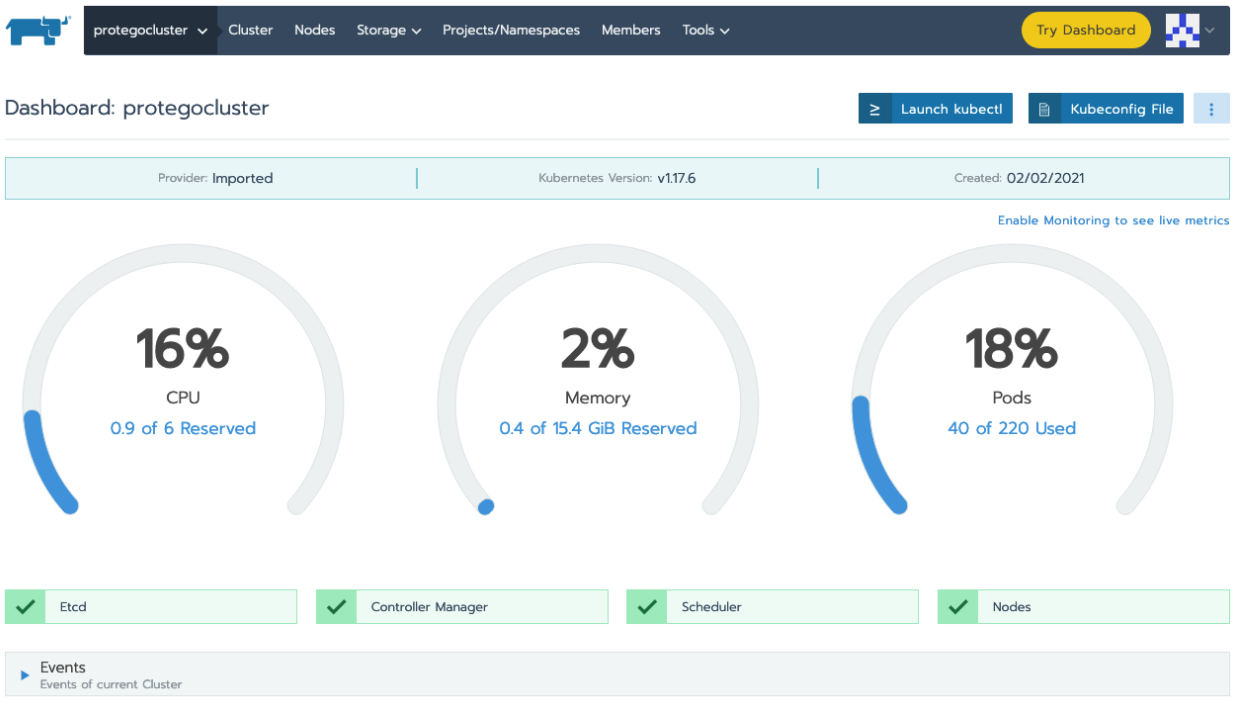


```
rancher@protegomaster:~$ kubectl get nodes
NAME              STATUS   ROLES               AGE    VERSION
192.168.249.232   Ready    controlplane,etcd   150d   v1.17.6
192.168.249.233   Ready    worker              150d   v1.17.6
192.168.249.234   Ready    worker              150d   v1.17.6
```

ID: S01.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The testers verified that Rancher prevents unauthorized users from logging in and creating a new cluster.

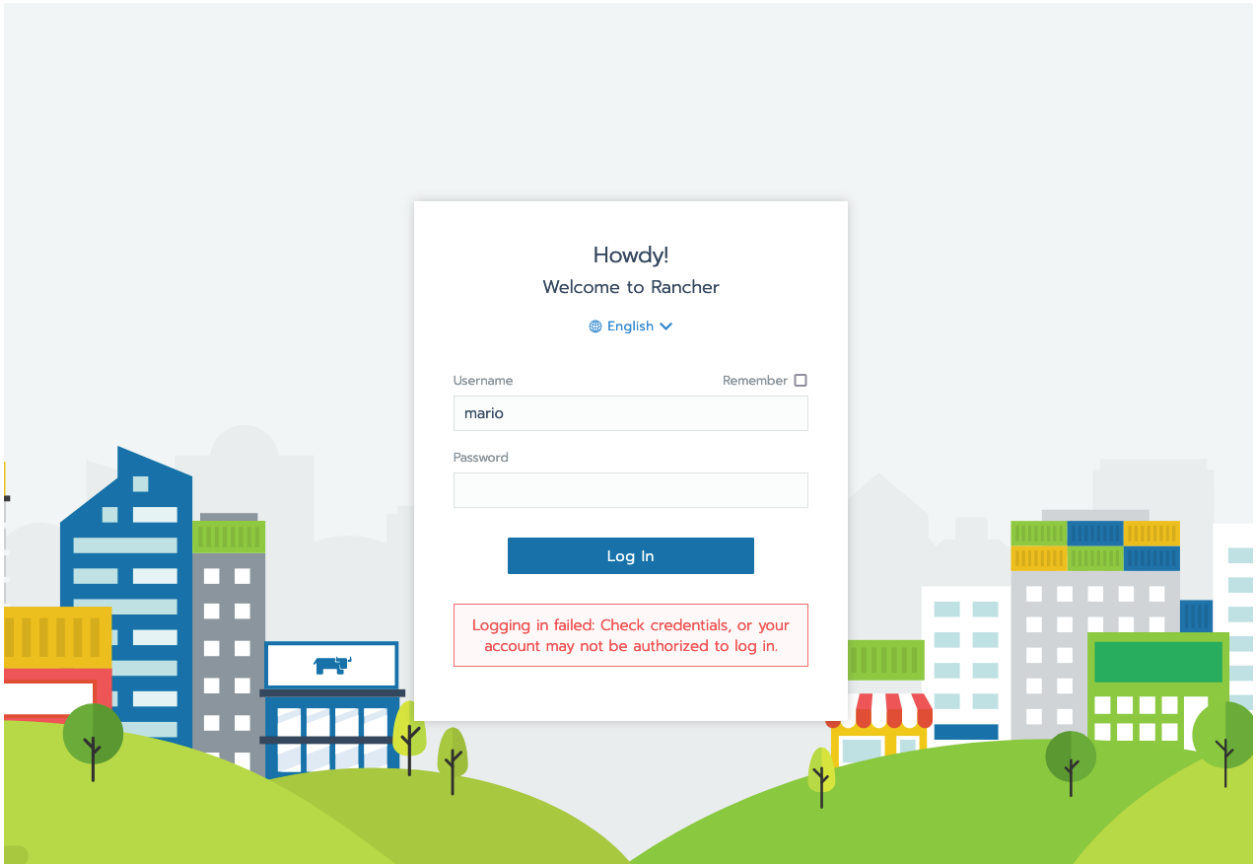

AT02: Install Data Gateway

ID: S02.1

Name: Operator installs Data Gateway

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to Rancher and installed the components of the Data Gateway, namely: the ELS-FHIR server, the KMS, and the Query Gateway. The successful deployment can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.

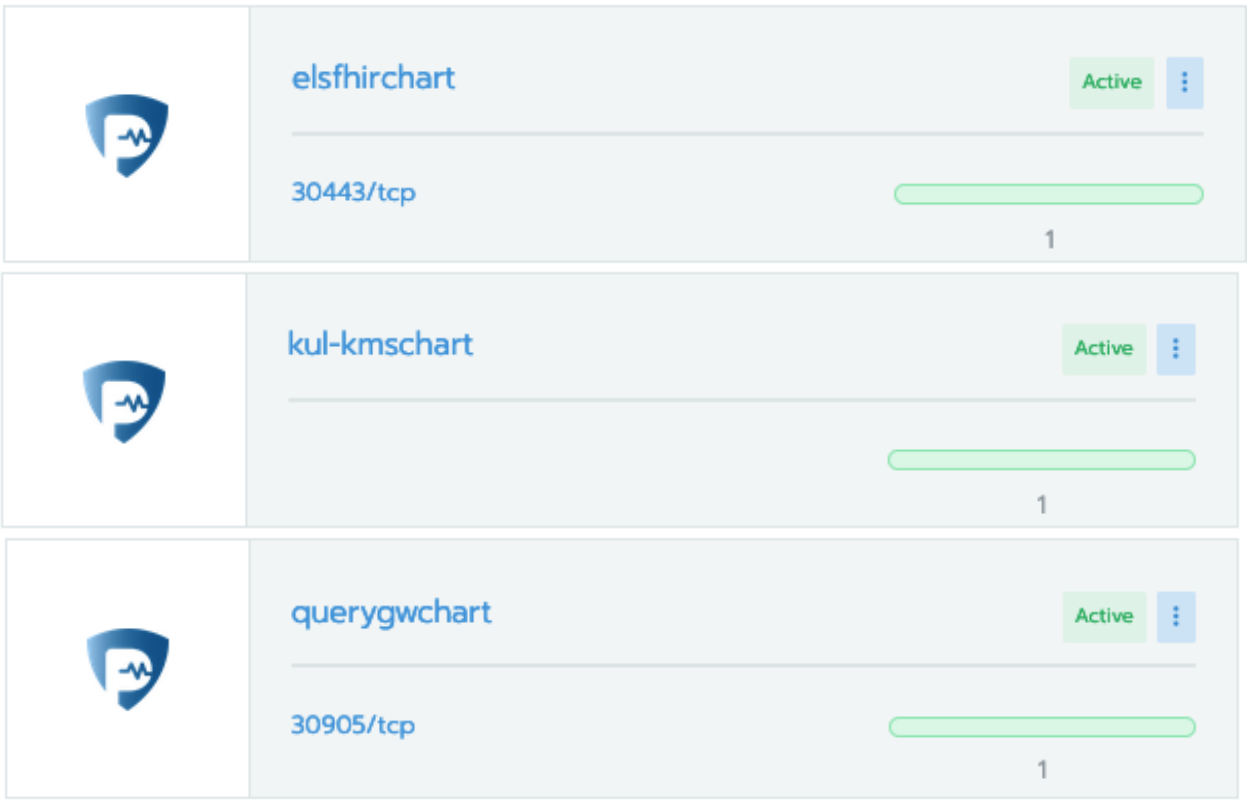


```
rancher@protegomaster:~$ kubectl get pods
NAME                          READY   STATUS    RESTARTS   AGE
els-fhir-6cf7f88c68-ghczf     1/1     Running   0          18d
kul-kms-6b564fd7c9-mqndd      1/1     Running   0          18d
querygw-84fc88cf65-gw8nn      1/1     Running   0          18d
```

ID: S02.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The testers verified that Rancher prevents unauthorized users from logging in and performing this task.

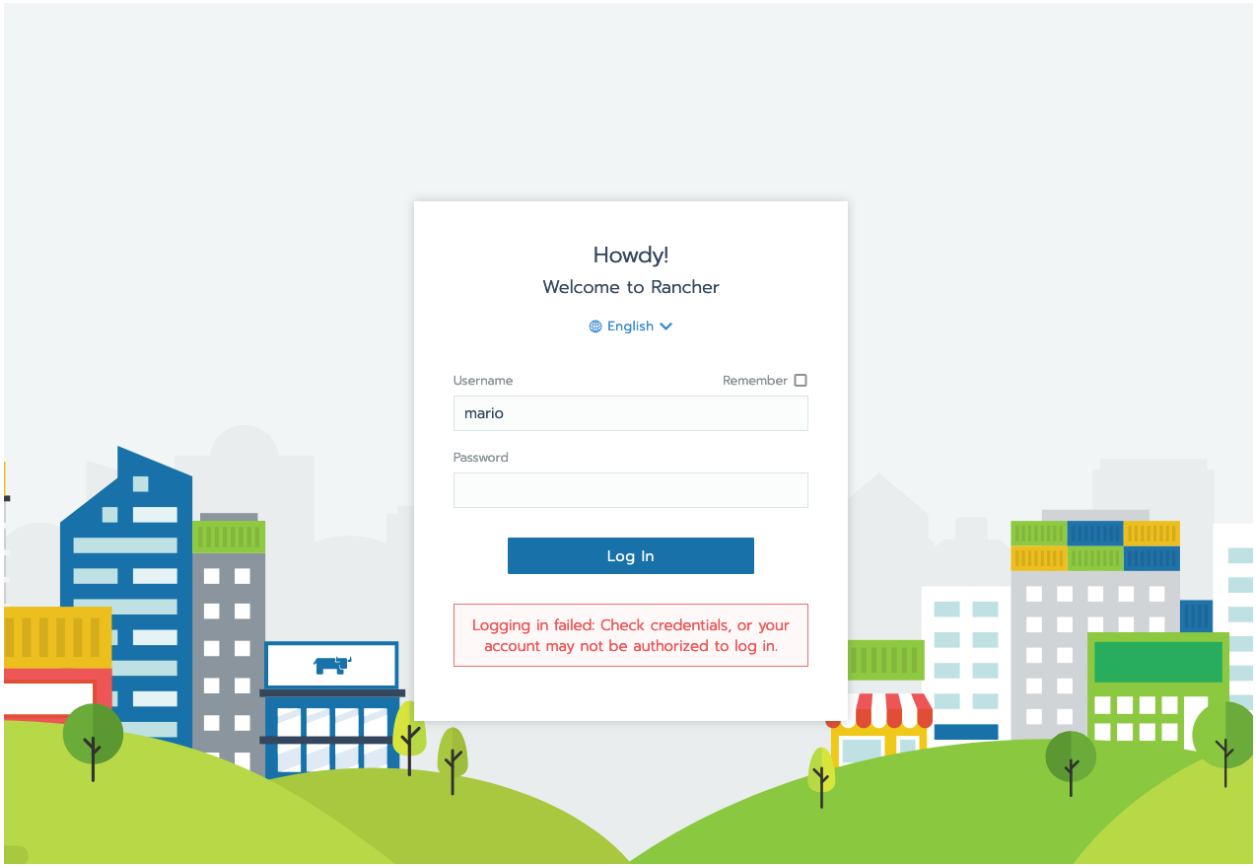

AT03: Install Network Slicing

ID: S03.1

Name: Operator installs Network Slicing

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to Rancher and installed the components of the Network Slicing, namely: the Runtime controller and the Ryu controller. The successful deployment can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.

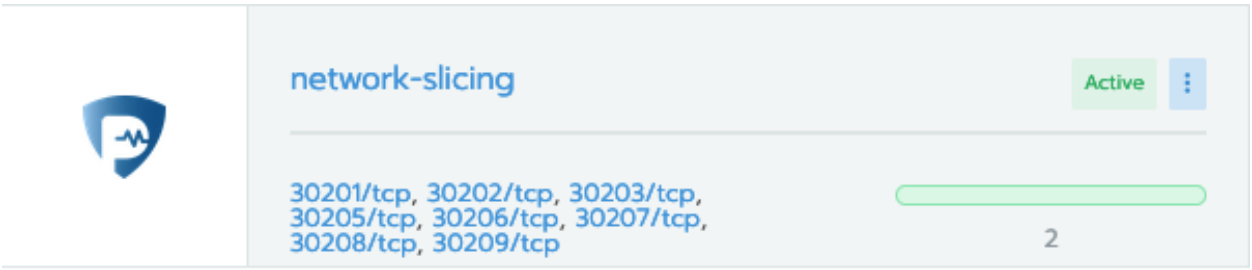


```
rancher@protegomaster:~$ kubectl get pods -n network-slicing
NAME                       READY   STATUS    RESTARTS   AGE
runtime-6c44bf9f68-vmdp2   1/1     Running   0          18d
ryu-57996798b-kbd8g        1/1     Running   0          18d
```

ID: S03.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: (see S02.2)

ID: S03.3

Name: Access point already registered

Priority: Medium

Evaluation: ✓ Passed

Description: The testers logged in to the controller web interface and tried to add a WTP (Wireless Termination Point) with the same MAC address of an already-existing WTP. As expected, the request was rejected.

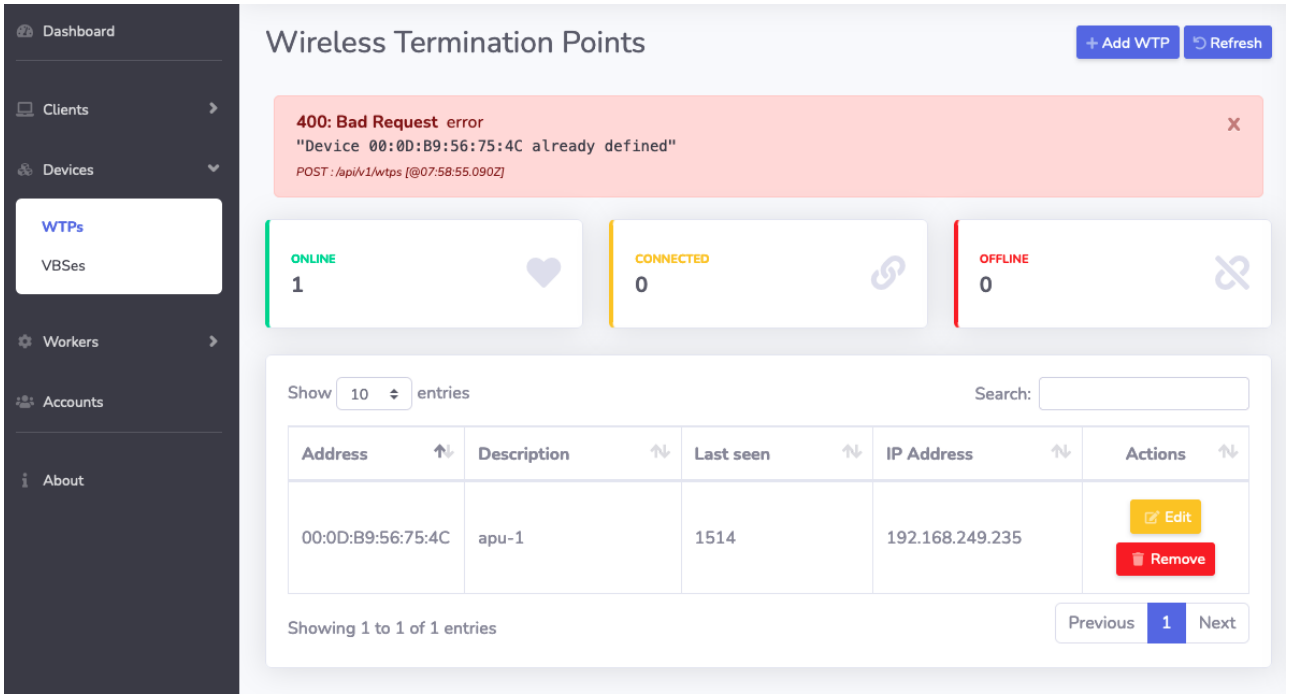

AT04: Install SIEM

ID: S04.1

Name: Operator installs SIEM

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to Rancher and installed the components of the SIEM. The successful deployment can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.




```
rancher@protegomaster:~$ kubectl get pods -n siem
NAME                          READY   STATUS    RESTARTS   AGE
analyzer-db497fcb4-8lx9p      1/1     Running   0          18d
elastic-5cb67cfcbb-xdldl      1/1     Running   0          18d
gsa-7958d5ddd4-xn69v          1/1     Running   0          18d
gvm-ff8f86cbc-dq4xd           1/1     Running   0          18d
kafka-59d6bc7558-hs69g        1/1     Running   1          18d
kibana-548fcf7fb9-sd7zk       1/1     Running   18         18d
logstash-c4f46774-6qcht       1/1     Running   0          18d
openvas-7c7c7b8c78-7brns      1/1     Running   0          18d
postgres-796654bf-mf6cw       1/1     Running   0          18d
scheduler-7754cf7cb9-4gbwq    1/1     Running   0          18d
wazuh-7d5f55dd9c-bhs44        1/1     Running   0          18d
zookeeper-8575598c7c-6h2fc    1/1     Running   0          18d
```

ID: S04.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: (see S02.2)

AT05: Install Continuous Authentication

ID: S05.1

Name: Operator installs Continuous Authentication

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to Rancher and installed the components of the Continuous Authentication, namely: two backends, the keystroke supervisor, and three DBMSs. The successful deployment can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.

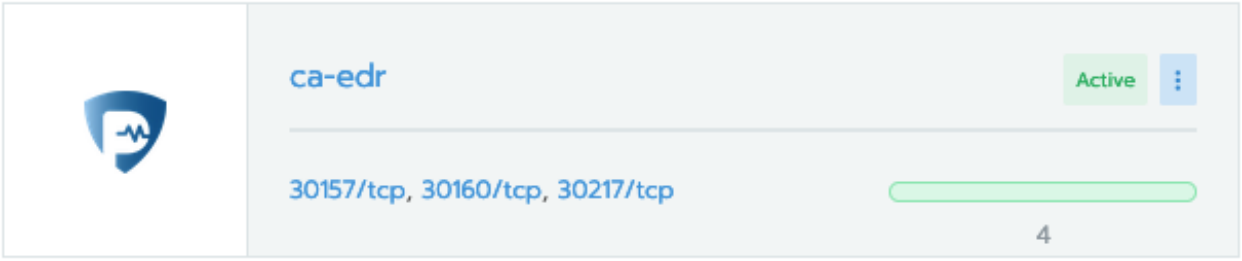

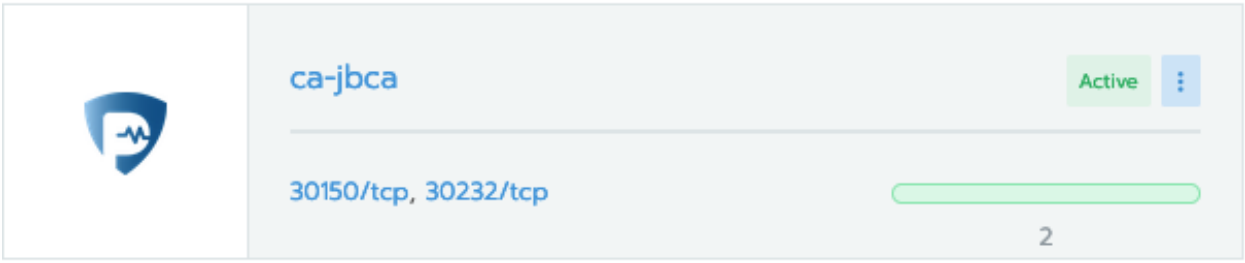


```
rancher@protegomaster:~$ kubectl get pods -n ca-edr
NAME                                  READY   STATUS    RESTARTS   AGE
api-c88d49858-x6xrq                   1/1     Running   0          18d
keystroke-supervisor-85db94b49-bz78p  1/1     Running   3          18d
mongo-0                               1/1     Running   0          18d
postgres-758484778b-qw9qk             1/1     Running   0          18d
```

```
rancher@protegomaster:~$ kubectl get pods -n ca-jbca
NAME                          READY   STATUS    RESTARTS   AGE
api-86444d6c85-rnt6c          1/1     Running   0          18d
postgres-6b57bf45bb-qjrjp     1/1     Running   0          18d
```

ID: S05.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: (see S02.2)

AT06: Install SSM

ID: S06.1

Name: Operator installs SSM

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to Rancher and installed the components of the SSM, namely: the backend, the adaptor, KeyCloak, the proxy, and two DBMSs. The successful deployment can be seen in the pictures below, both on the Rancher GUI and on the CLI via `kubectl`.

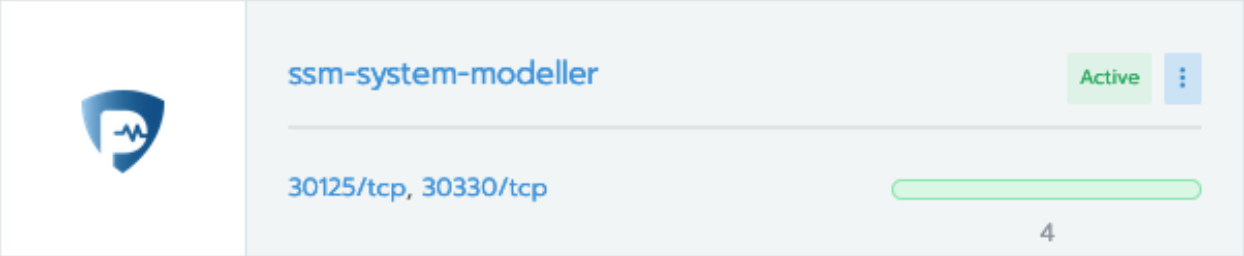

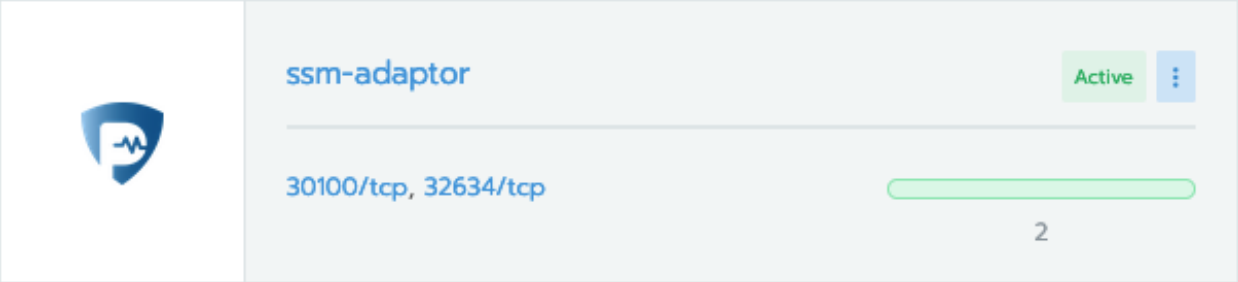


```
rancher@protegomaster:~$ kubectl get pods -n ssm-system-modeller
NAME                                  READY   STATUS    RESTARTS   AGE
keycloak-deployment-dc8fc885b-gzzps   1/1     Running   0          18d
mongo-deployment-575bbcf49c-lxqwm     1/1     Running   0          18d
proxy-deployment-7df7cf855b-6tv6l     1/1     Running   0          18d
ssm-deployment-dddc7f45f-vzdpn        1/1     Running   0          18d
```

```
rancher@protegomaster:~$ kubectl get pods -n ssm-adaptor
NAME                            READY   STATUS    RESTARTS   AGE
mongodb-5f94ccbf6d-s55qj        1/1     Running   0          18d
ssm-adaptor-7b78cb9798-5gf5k    1/1     Running   0          18d
```

ID: S06.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: (see S02.2)

AT07: Register user

ID: S07.1

Name: Administrator registers a new user

Priority: High

Evaluation: ✓ Passed

Description: There is not a central system in ProTego to register users. However, all the ProTego components have their own functionality to register users (the image shows the SIEM log in process).

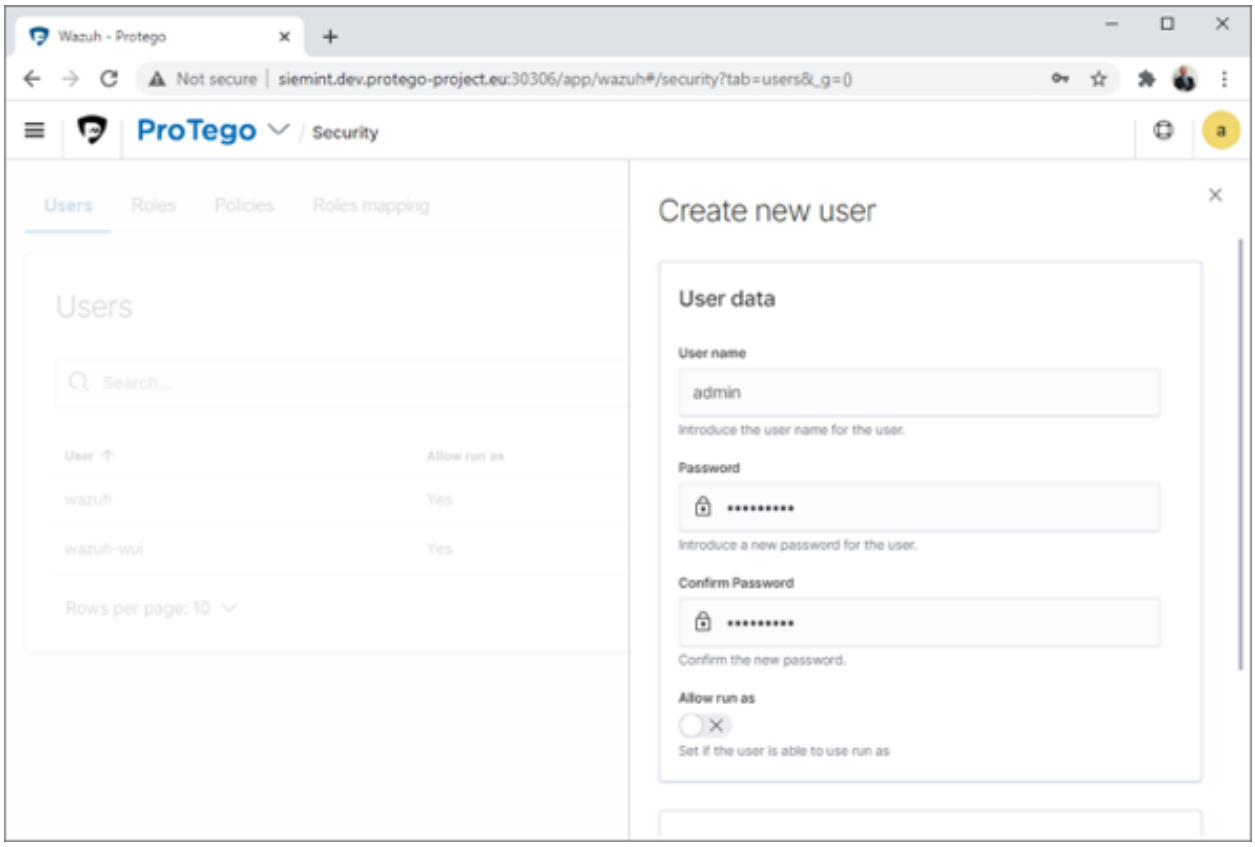

ID: S07.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The testers verified that each component prevents unauthorized users from performing this task.

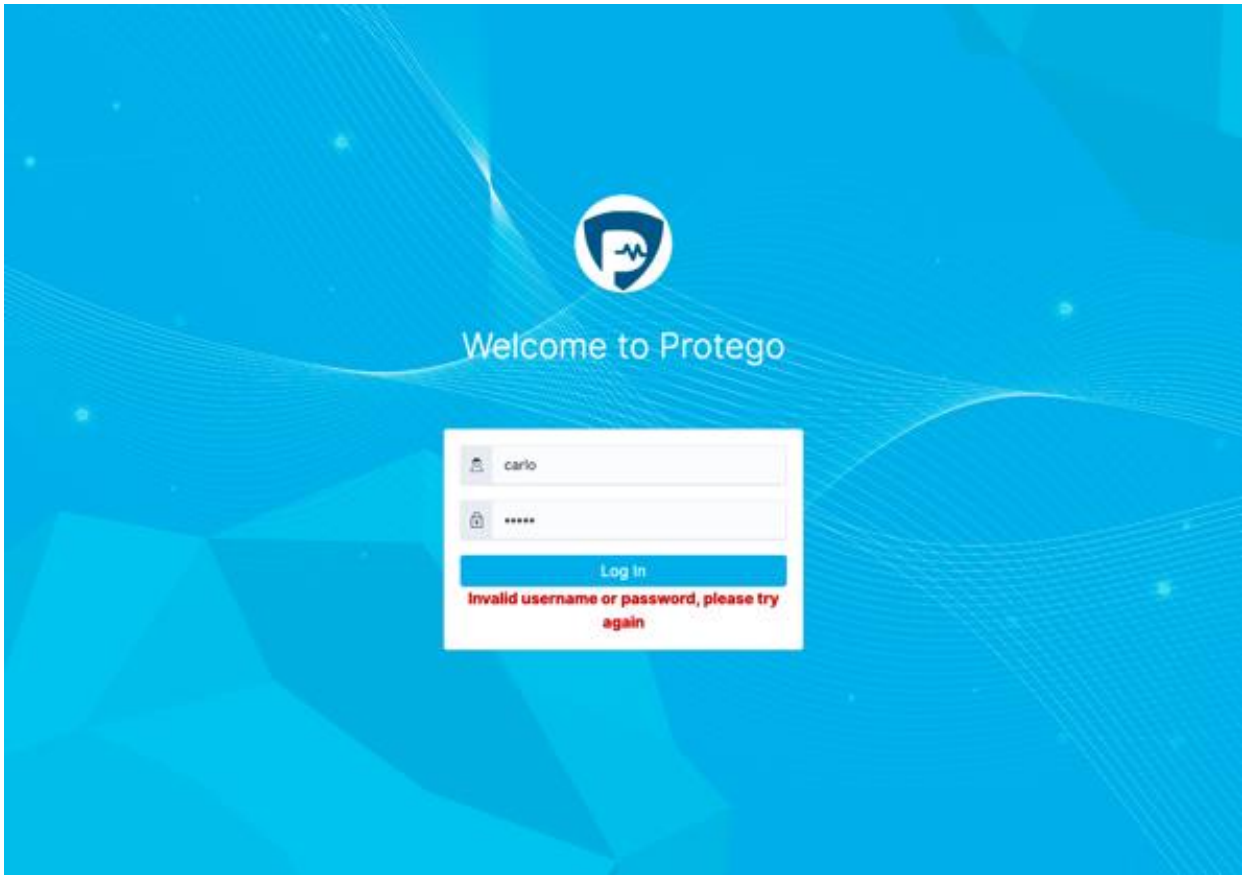

ID: S07.3

Name: User identifier already taken

Priority: Medium

Evaluation: ✓ Passed

Description: The testers verified that each component prevents registration of users with same ID.

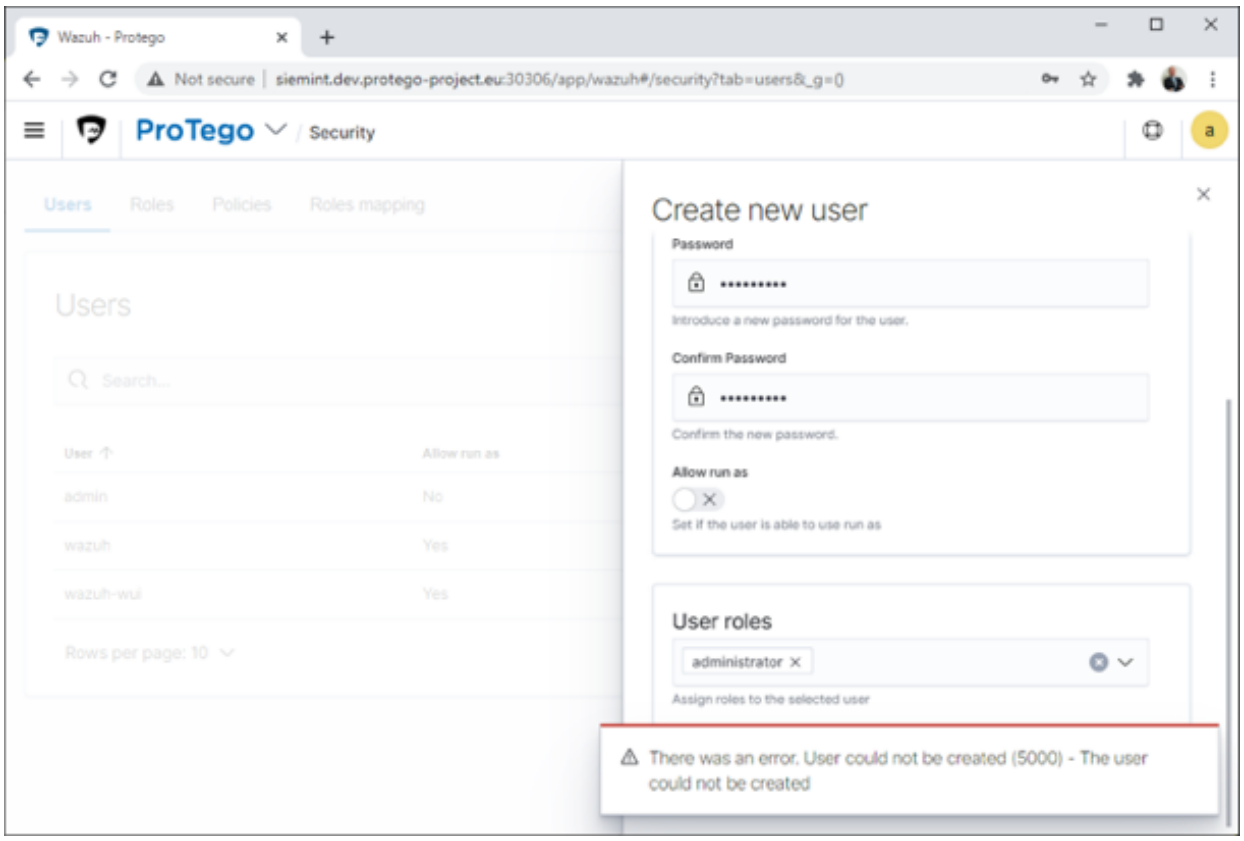

AT08: Conduct first-time risk assessment

ID: S08.1

Name: Security operator assesses risks

Priority: High

Evaluation: ✓ Passed

Description: The testers logged in to the SSM, then they modeled the OSR infrastructure (including the FoodCoach application) and calculated the associated risks.

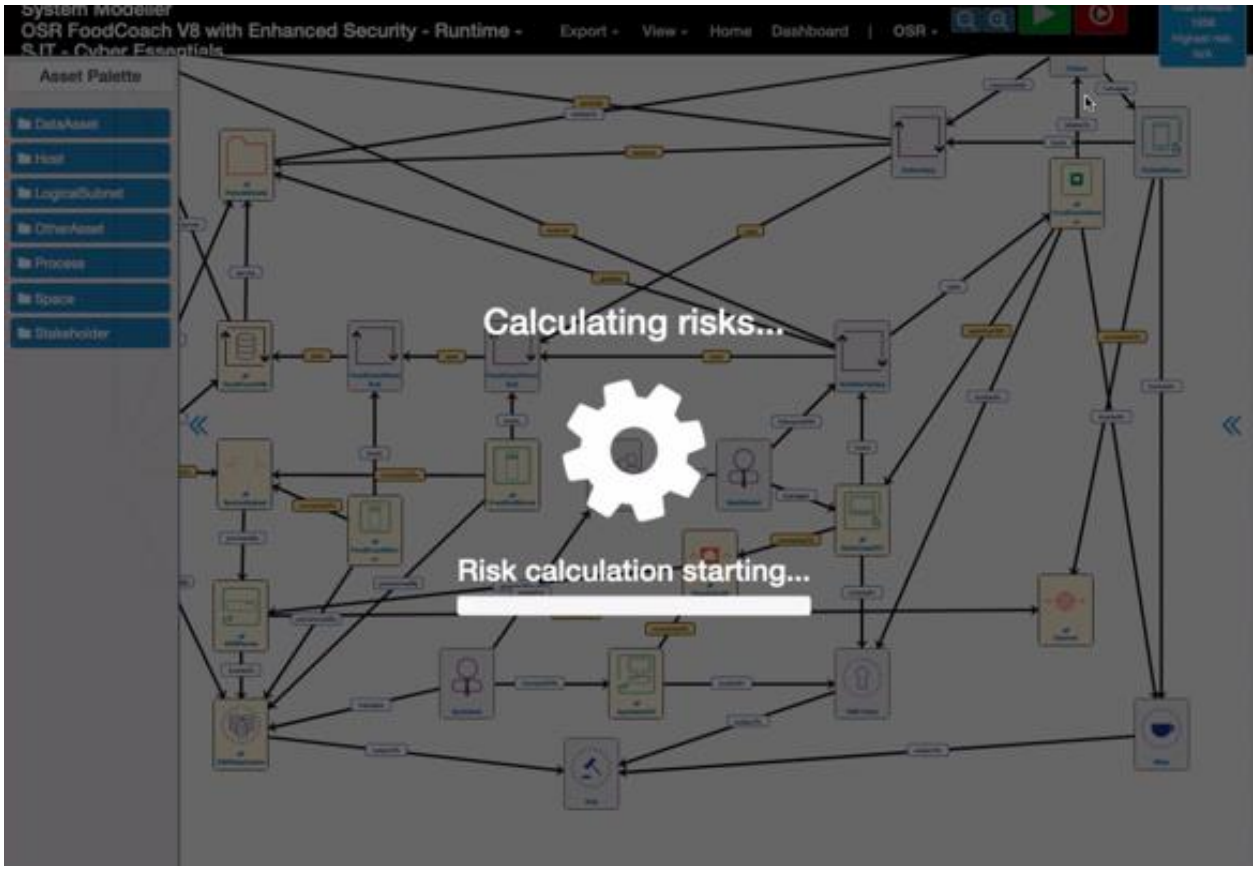

ID: S08.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The tester tried to log in to the SSM with credentials that were not associated with a Security operator, and the operation failed.

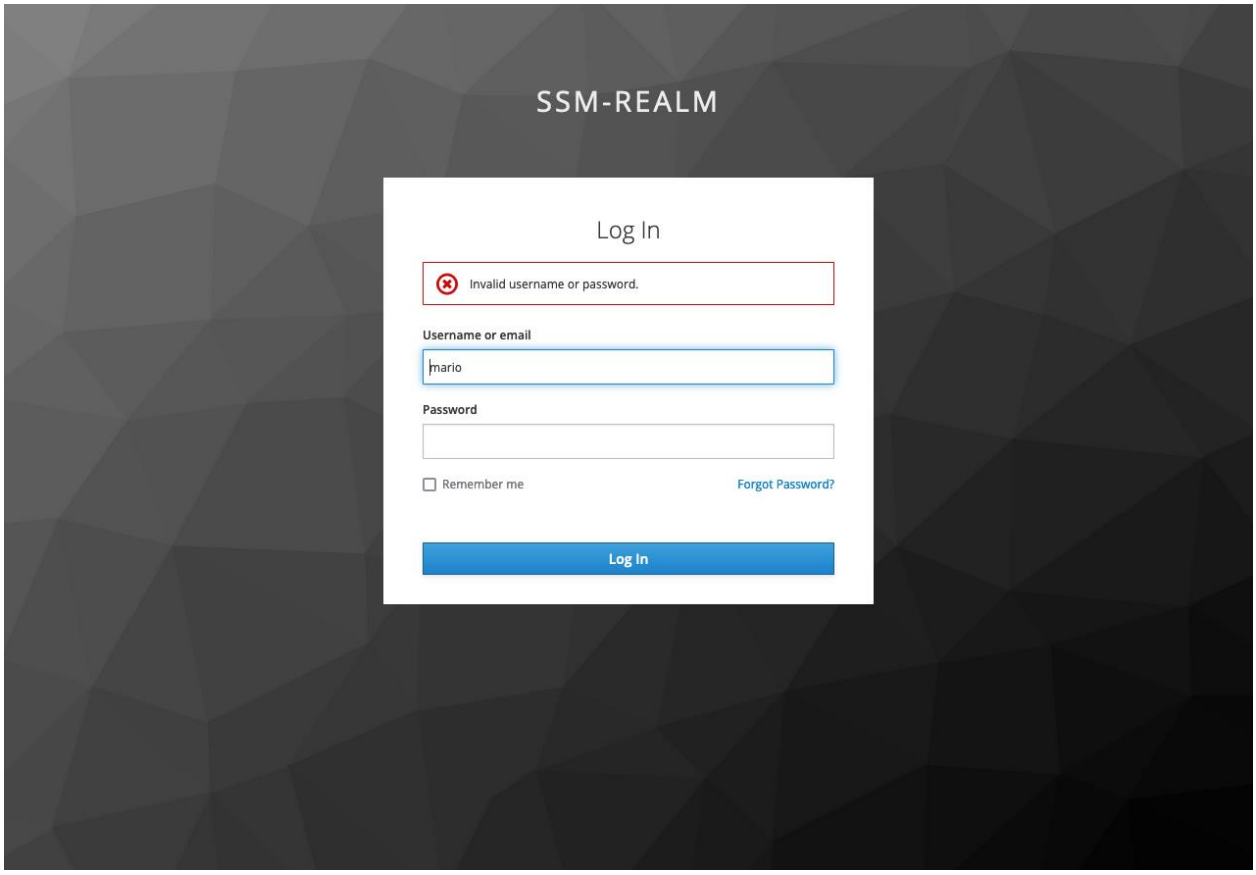

ID: S08.3

Name: First-time assessment already conducted

Priority: Medium

Evaluation: ✗ Failed

Description: The testers verified that no check is currently done on the unicity of the name of the models, thus allowing a user to potentially create multiple models with the same name. This aspect is currently under development to support commercialization of the SSM. More specifically, this check is part of a wider plan for risk model management within the SSM tool, including copying and version control of models, and multiple users accessing the same SSM instance and sharing risk model files. Users will have their own workspace with model files in and will be able to share them with other users (which complicates the issue of model naming).

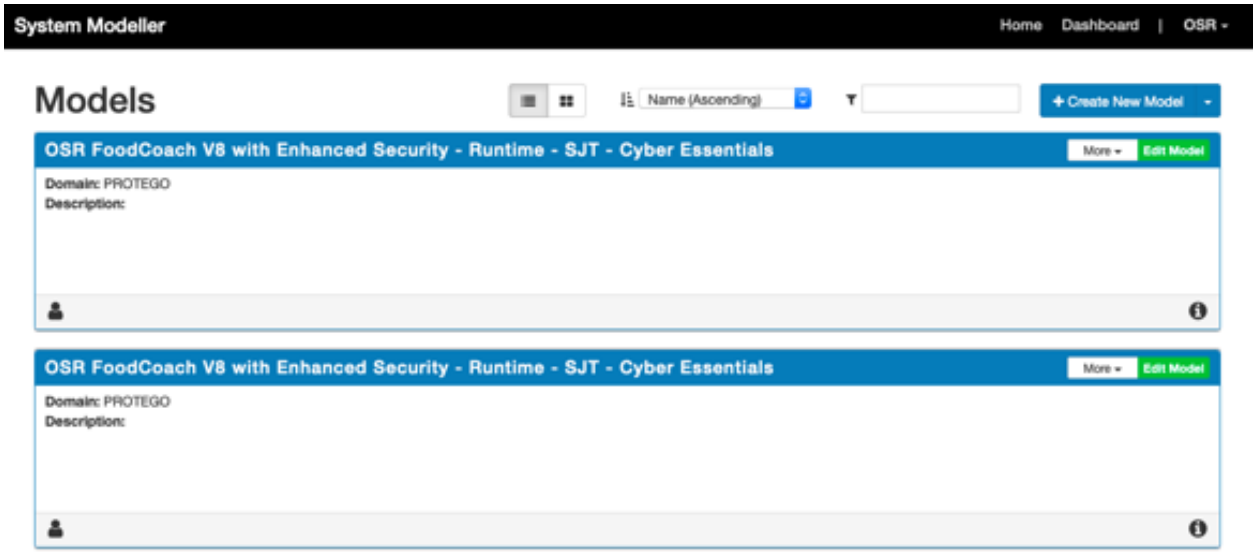

AT09: Assess prospective risks to the infrastructure

ID: S09.1

Name: Security operator assesses application risks

Priority: High

Evaluation: ✓ Passed

Description: The testers selected the application asset modelled in the infrastructure, and the SSM displayed the associated risks.

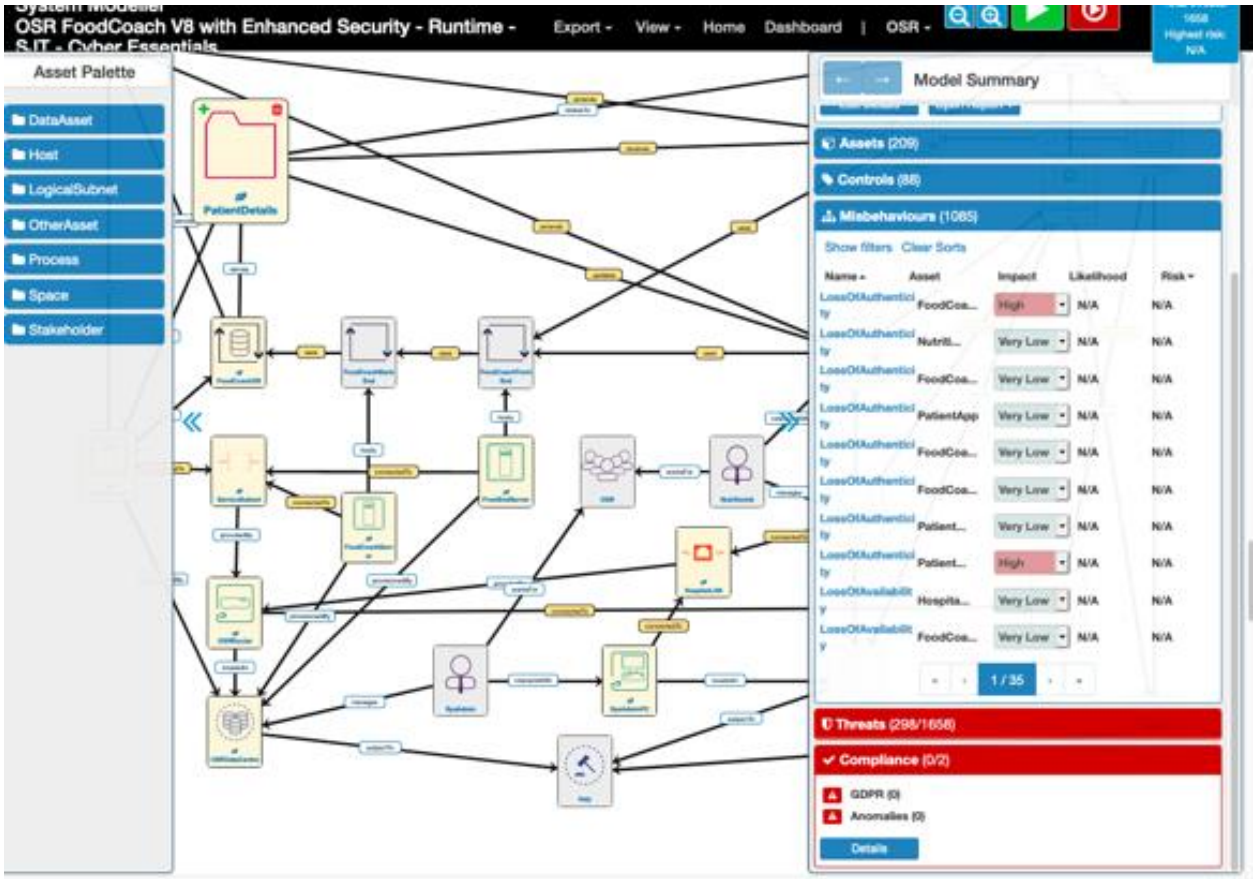

ID: S09.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The tester tried to log in to the SSM with credentials that were not associated with a Security operator, and the operation failed.

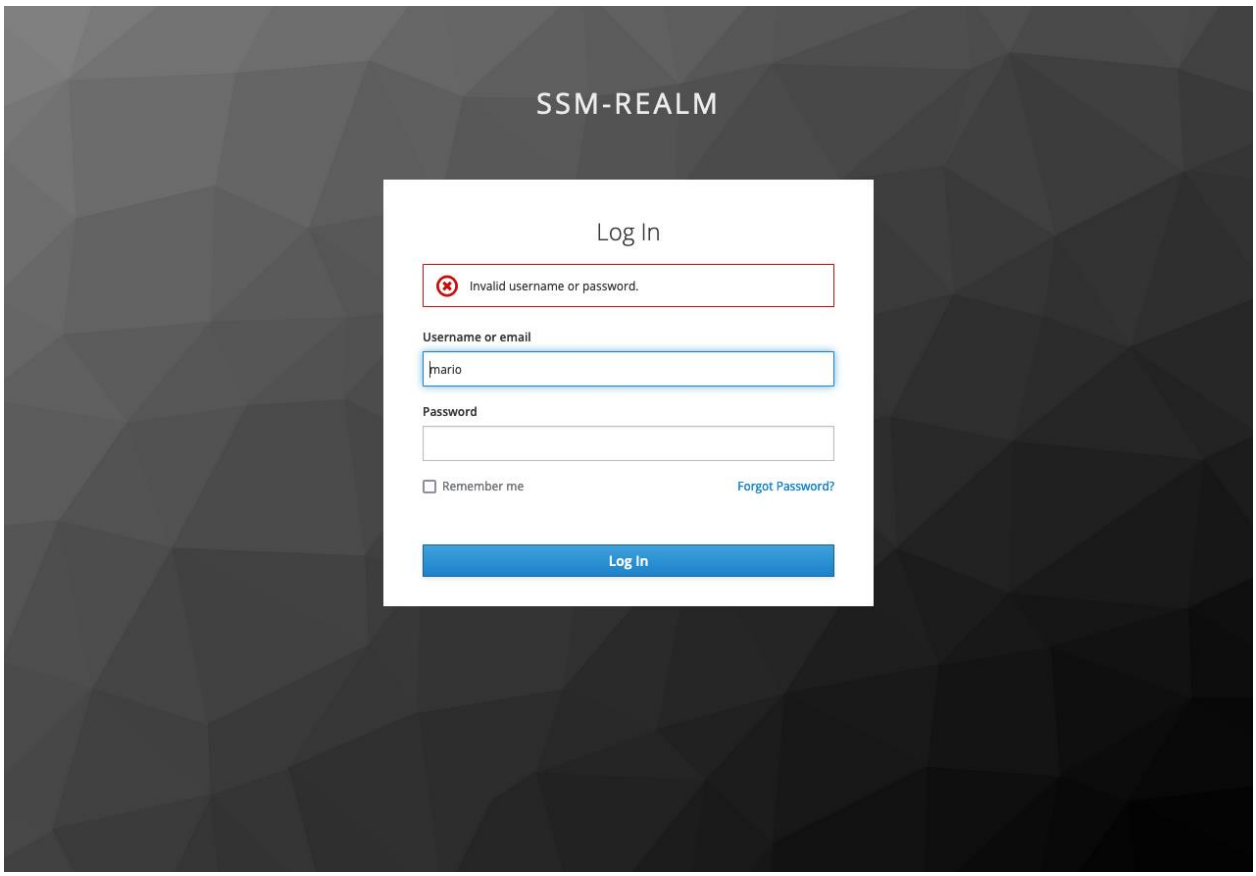

ID: S09.3

Name: First-time risk assessment is still to be conducted

Priority: Medium

Evaluation: ✓ Passed

Description: the testers are made aware that the risk assessment was not yet performed, and no models are available.



AT10: Install application

ID: S10.1

Name: System operator installs a new application

ID: S10.2

Name: Insufficient authorization

ID: S10.3

Name: Already existing application

Evaluation: Outdated

Description: This Acceptance Test has been marked as ***Outdated***. The reason for this is that the ProTego toolkit had originally been planned to manage all the components and the applications in a centralized way. That way, the toolkit would have a list of the applications installed so far and would be able to reject the installation of an already-installed application. Though, the benefit of this would not be particularly significant. On the other hand, installing the application independently from the toolkit offers much more flexibility, which is what is currently being done at both hospitals, with FoodCoach and Pocket EHR. Within this frame of reference, all the three scenarios actually pass, but not in a way that is strictly controlled by the ProTego toolkit. In fact, an operator is able to install the application, but to do so, they must have authorization to access the infrastructure and the application must not be already installed.

AT11: Configure application network slices

ID: S11.1
Name: Security operator configures network slices
Priority: High
Evaluation: ✓ Passed
Description: The tester logged in to the Network Slicing controller, and then created two Wi-Fi slices. Most importantly, the Quantum of each slice was configured, which is used to determine the rate of the traffic passing through that slice.

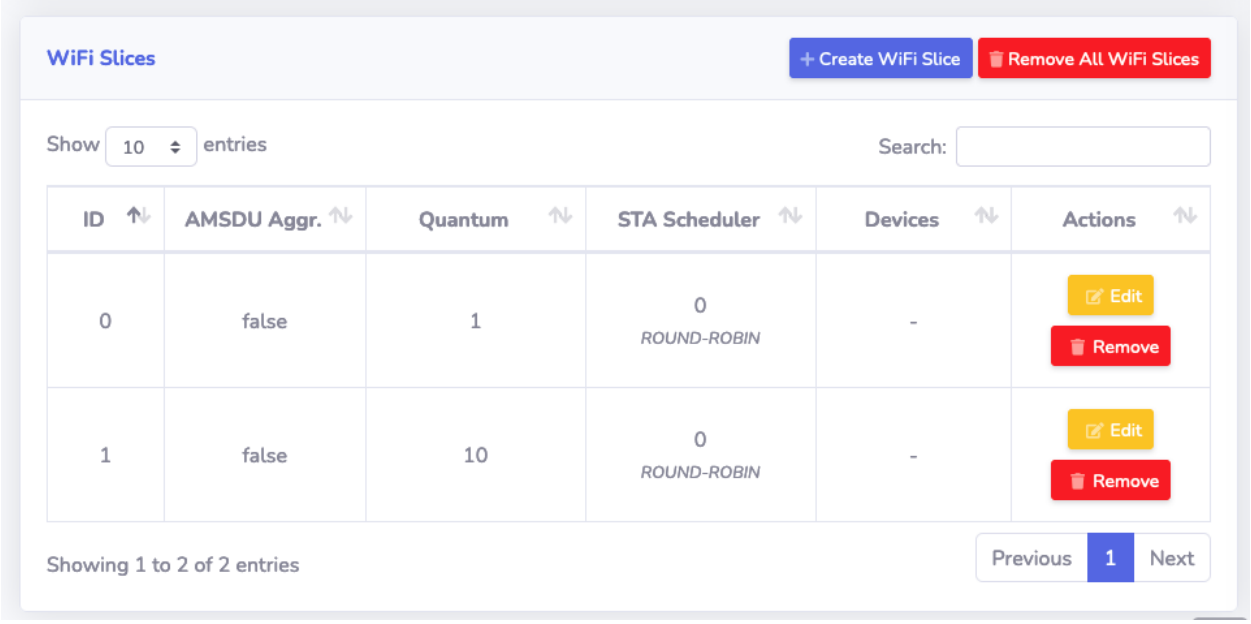

ID: S11.2
Name: Insufficient authorization
Priority: Medium
Evaluation: ✓ Passed
Description: The tester tried to log in to the Network Slicing web interface with credentials that were not associated with a Network operator, and the operation failed.

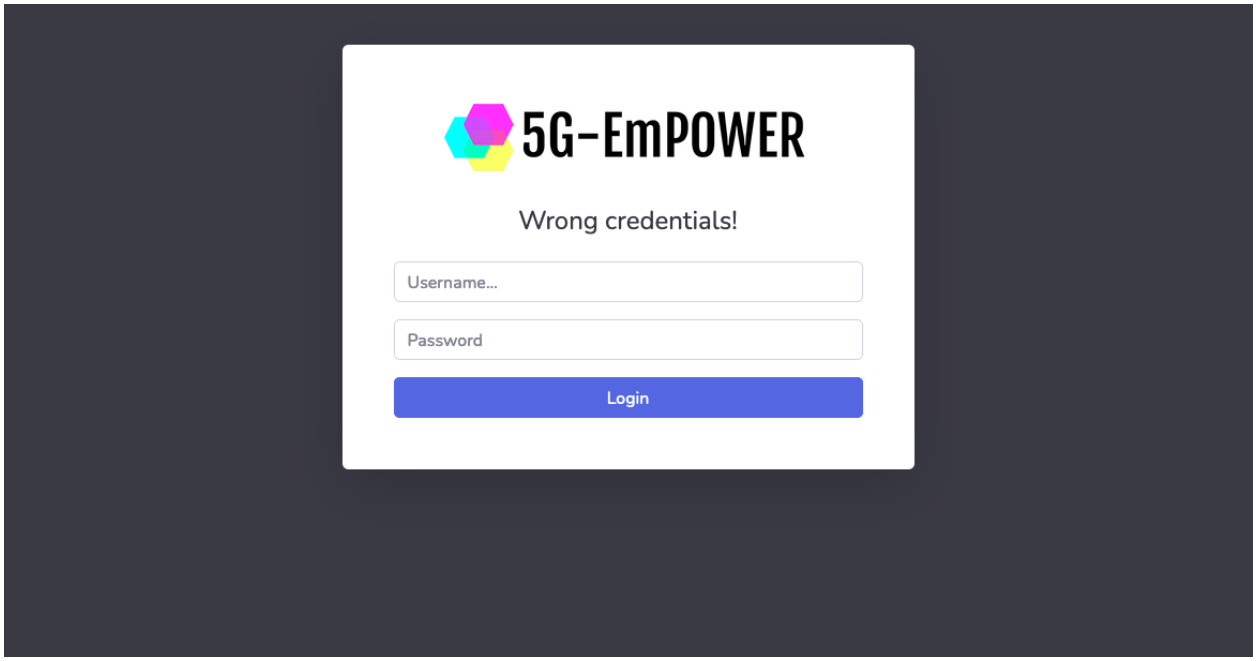

ID: S11.3
Name: Network slices are already configured
Priority: Medium
Evaluation: ✓ Passed
Description: The application network slice was already configured. The testers tried to configure it again by using the same ID, and the operation failed as expected.

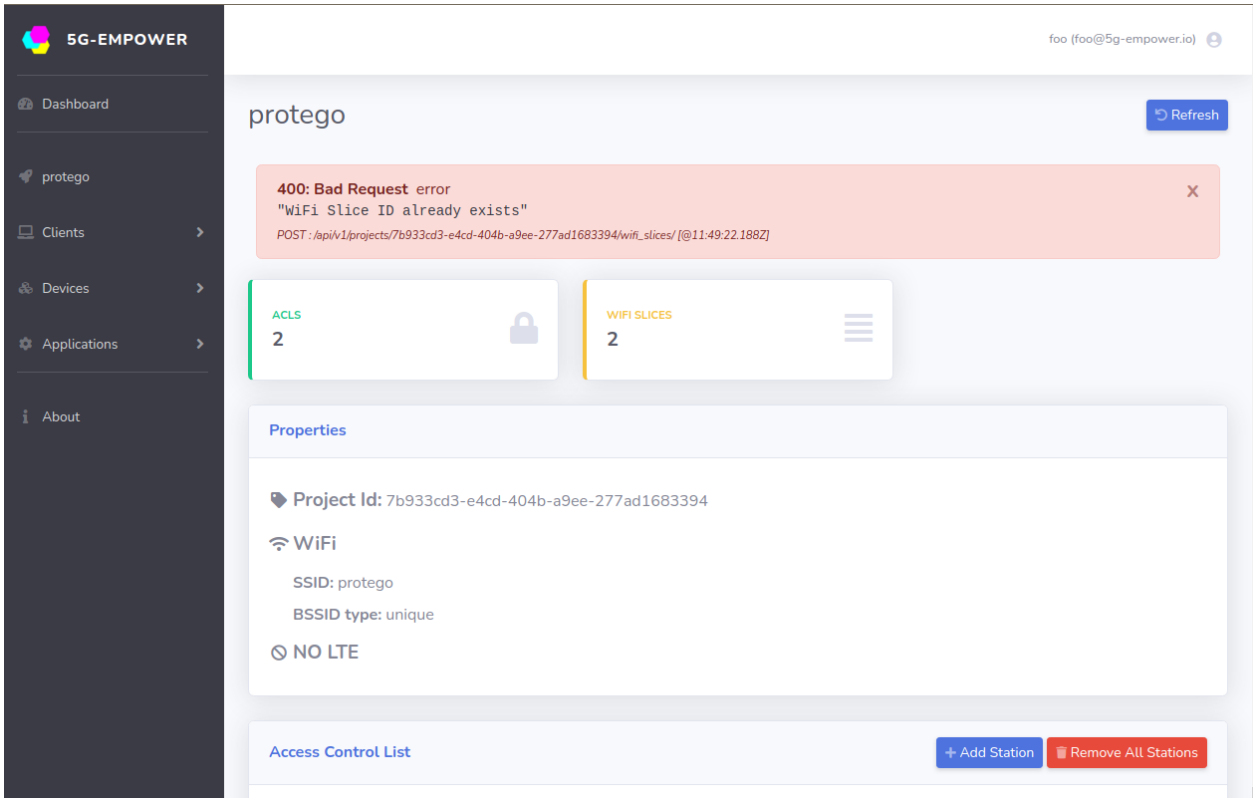

AT12: Configure application logging mechanism

ID: S12.1

Name: Security operator specifies logging mechanism

Priority: High

Evaluation: ✓ Passed

Description: The testers configured the logging mechanism to monitor the application, by performing three steps. Firstly, the SIEM agent was configured so that it could read the logs of the application and so that it could be registered in the SIEM (as shown in screenshot 1). Then, a decoder was created (screenshot 2) and it was tailored specifically to read the application logs. Lastly, rules were defined (screenshot 3) to detect specific events from the information that was extracted by the decoder.

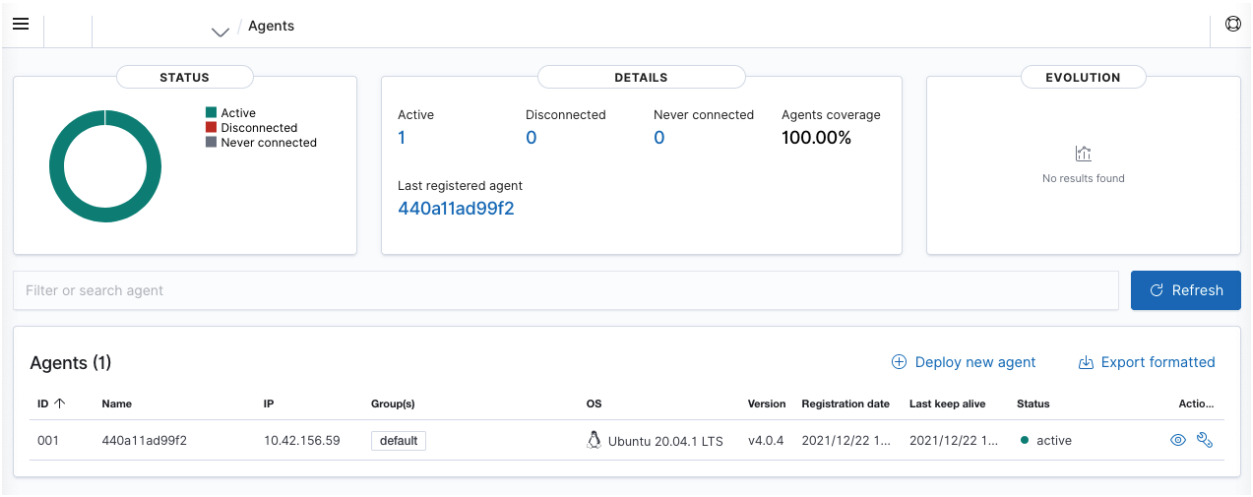

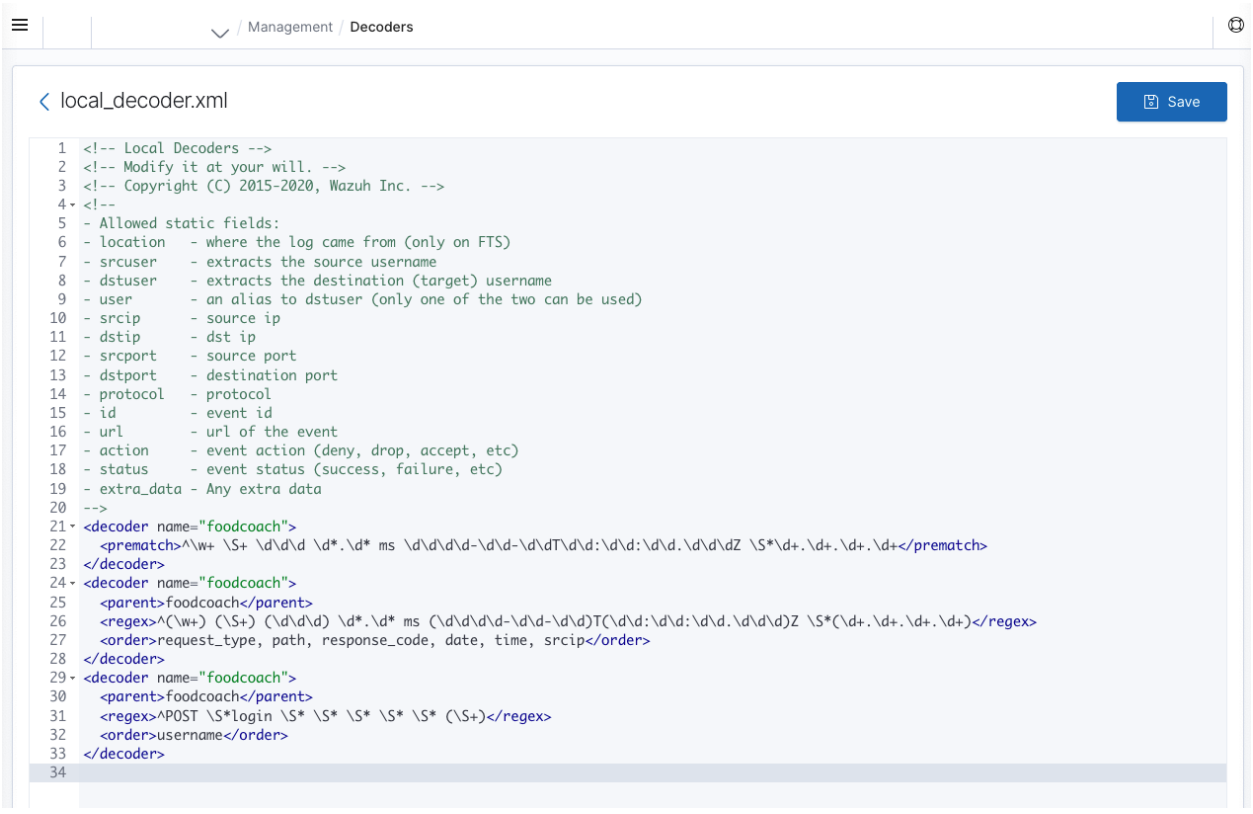

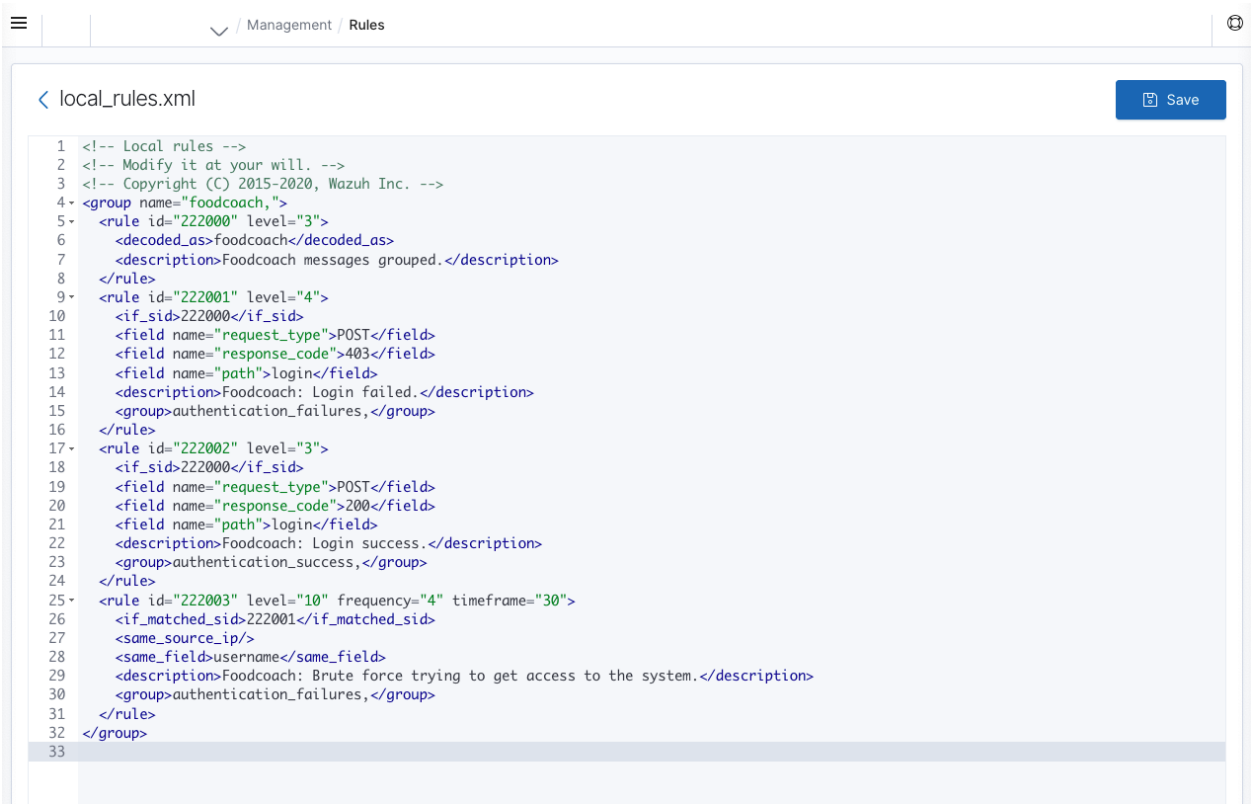
Management / Rules
local_rules.xml
Save
<!-- Local rules -->
<!-- Modify it at your will. -->
<!-- Copyright (C) 2015-2020, Wazuh Inc. -->
<group name="foodcoach,">
<rule id="222000" level="3">
<decoded_as>foodcoach</decoded_as>
<description>Foodcoach messages grouped.</description>
</rule>
<rule id="222001" level="4">
<if_sid>222000</if_sid>
<field name="request_type">POST</field>
<field name="response_code">403</field>
<field name="path">login</field>
<description>Foodcoach: Login failed.</description>
<group>authentication_failures,</group>
</rule>
<rule id="222002" level="3">
<if_sid>222000</if_sid>
<field name="request_type">POST</field>
<field name="response_code">200</field>
<field name="path">login</field>
<description>Foodcoach: Login success.</description>
<group>authentication_success,</group>
</rule>
<rule id="222003" level="10" frequency="4" timeframe="30">
<if_matched_sid>222001</if_matched_sid>
<same_source_ip/>
<same_field>username</same_field>
<description>Foodcoach: Brute force trying to get access to the system.</description>
<group>authentication_failures,</group>
</rule>
</group>

ID: S12.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The testers tried to log in SIEM tool with credential not associated with Network operator and the operation failed.

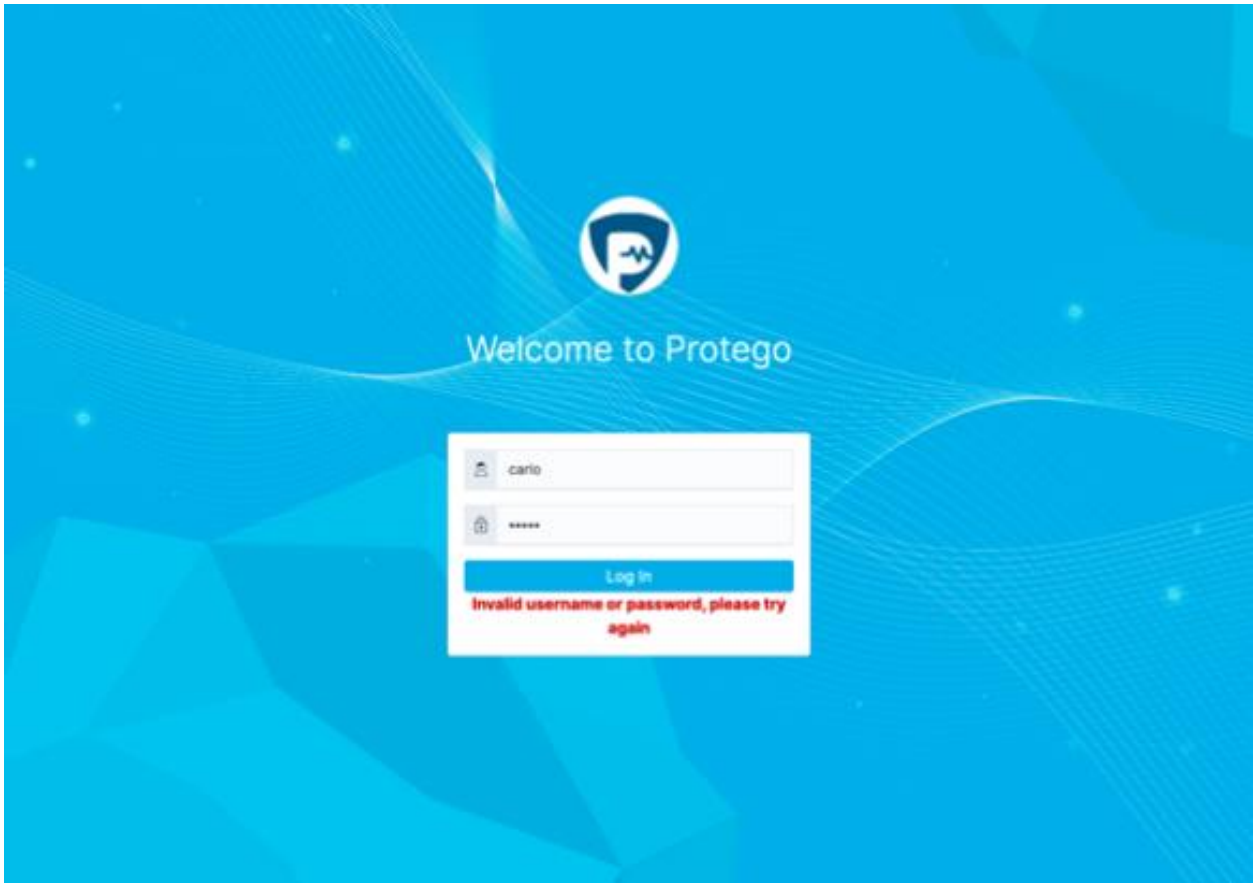

AT13: Specify application access control

ID: S13.1

Name: Security operator specifies access control

Priority: High

Evaluation: ✓ Passed

Description: There is not a centralized system in ProTego to specify access control rules. However, the testers verify the possibility to specify the application access control by manually intervening on the policy.

There is a version of the access control component (v6.1) that is not integrated into the demo and is only implemented in WP5. This version is based on OPA, allows better configuration of access control policies, and offers a more user-friendly way to configure the access control.

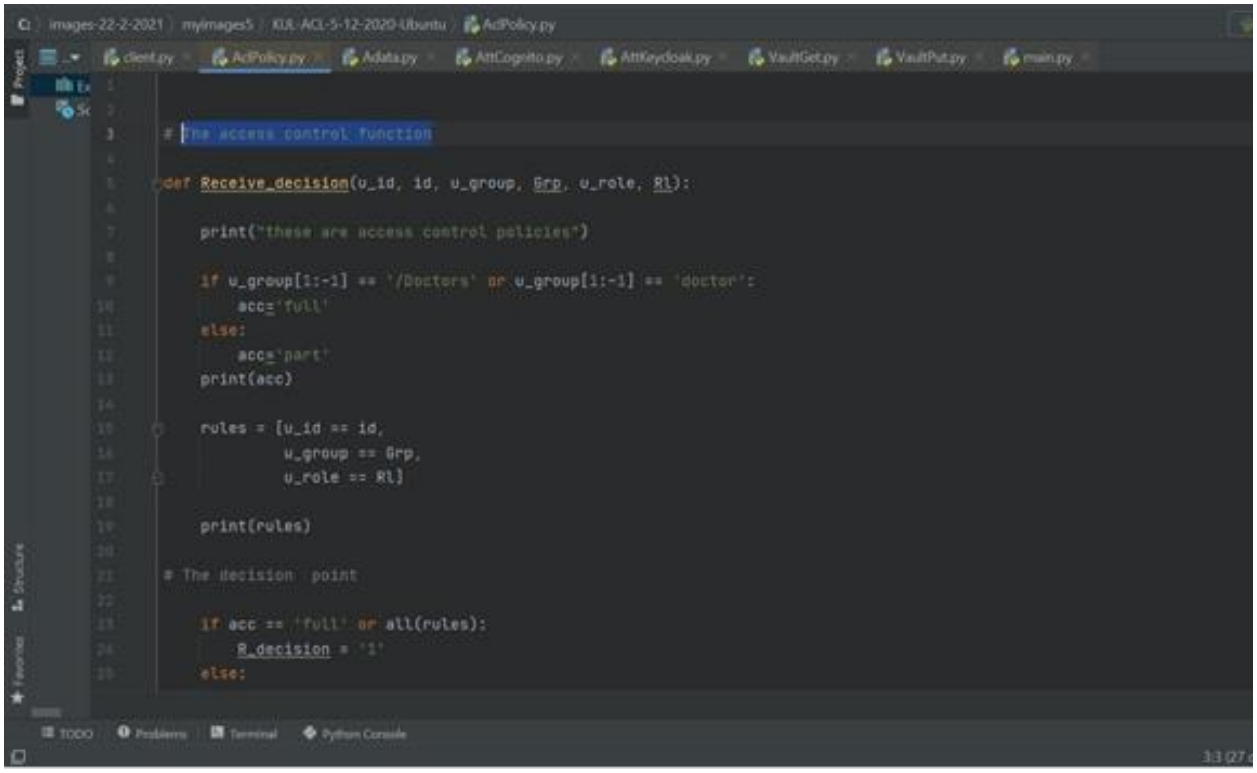

```
# The access control function

def Receive_decision(u_id, id, u_group, Grp, u_role, Rl):

    print("these are access control policies")

    if u_group[1:-1] == '/Doctors' or u_group[1:-1] == 'doctor':
        acc='full'
    else:
        acc='part'
    print(acc)

    rules = [u_id == id,
             u_group == Grp,
             u_role == Rl]

    print(rules)

# The decision  point

    if acc == 'full' or all(rules):
        R_decision = '1'
    else:
```

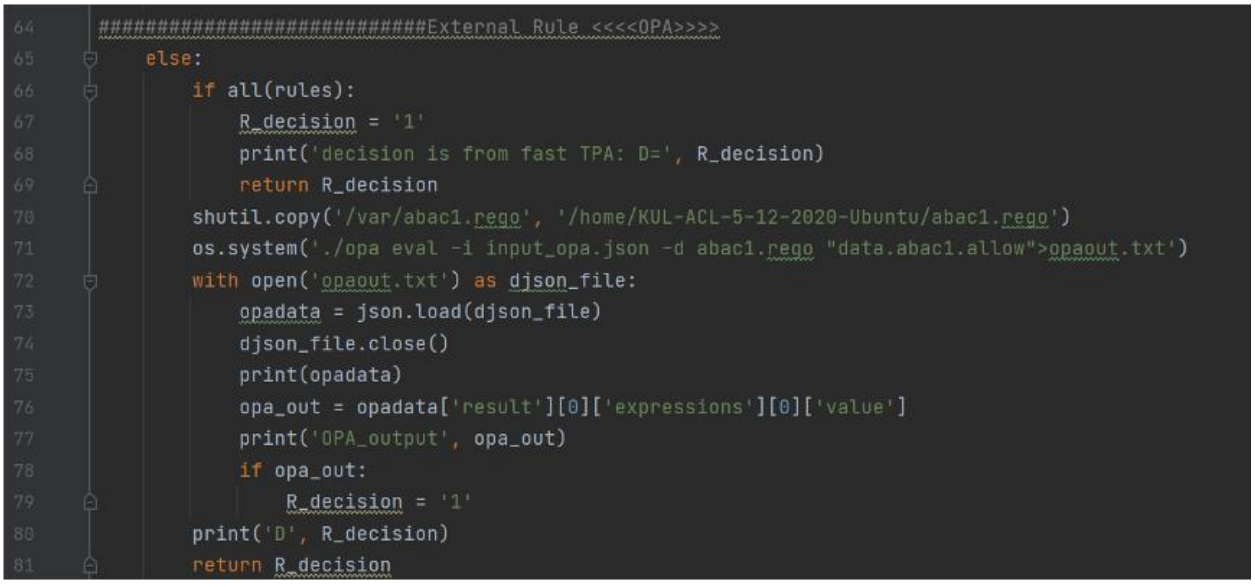

```
##############################External Rule <<<<OPA>>>>
    else:
        if all(rules):
            R_decision = '1'
            print('decision is from fast TPA: D=', R_decision)
            return R_decision
        shutil.copy('/var/abac1.rego', '/home/KUL-ACL-5-12-2020-Ubuntu/abac1.rego')
        os.system('./opa eval -i input_opa.json -d abac1.rego "data.abac1.allow">opaout.txt')
        with open('opaout.txt') as djson_file:
            opadata = json.load(djson_file)
            djson_file.close()
            print(opadata)
            opa_out = opadata['result'][0]['expressions'][0]['value']
            print('OPA_output', opa_out)
            if opa_out:
                R_decision = '1'
        print('D', R_decision)
        return R_decision
```

ID: S13.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The testers verify that only authorized user can specify the access control policy (i.e., change or update the security policies in place).

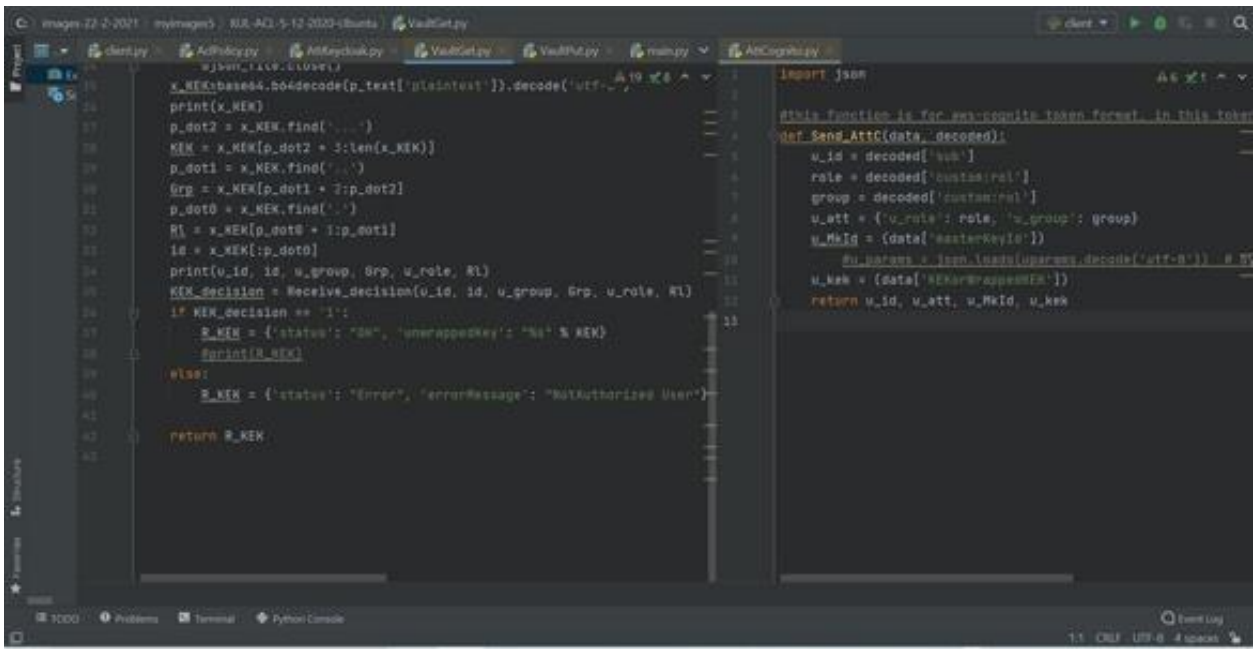

AT14: Configure mobile device

ID: S14.1

Name: Security Operator configures continuous authentication agent

Priority: High

Evaluation: ✓ Passed

Description: The mobile device can be configured by performing three steps. First, two applications need to be installed: the EDR and the FoodCoach application (or Pocket EHR). Then, the proper keyboard must be selected. Lastly, admin permissions must be granted, which allow to lock the screen when suspicious activity is detected.

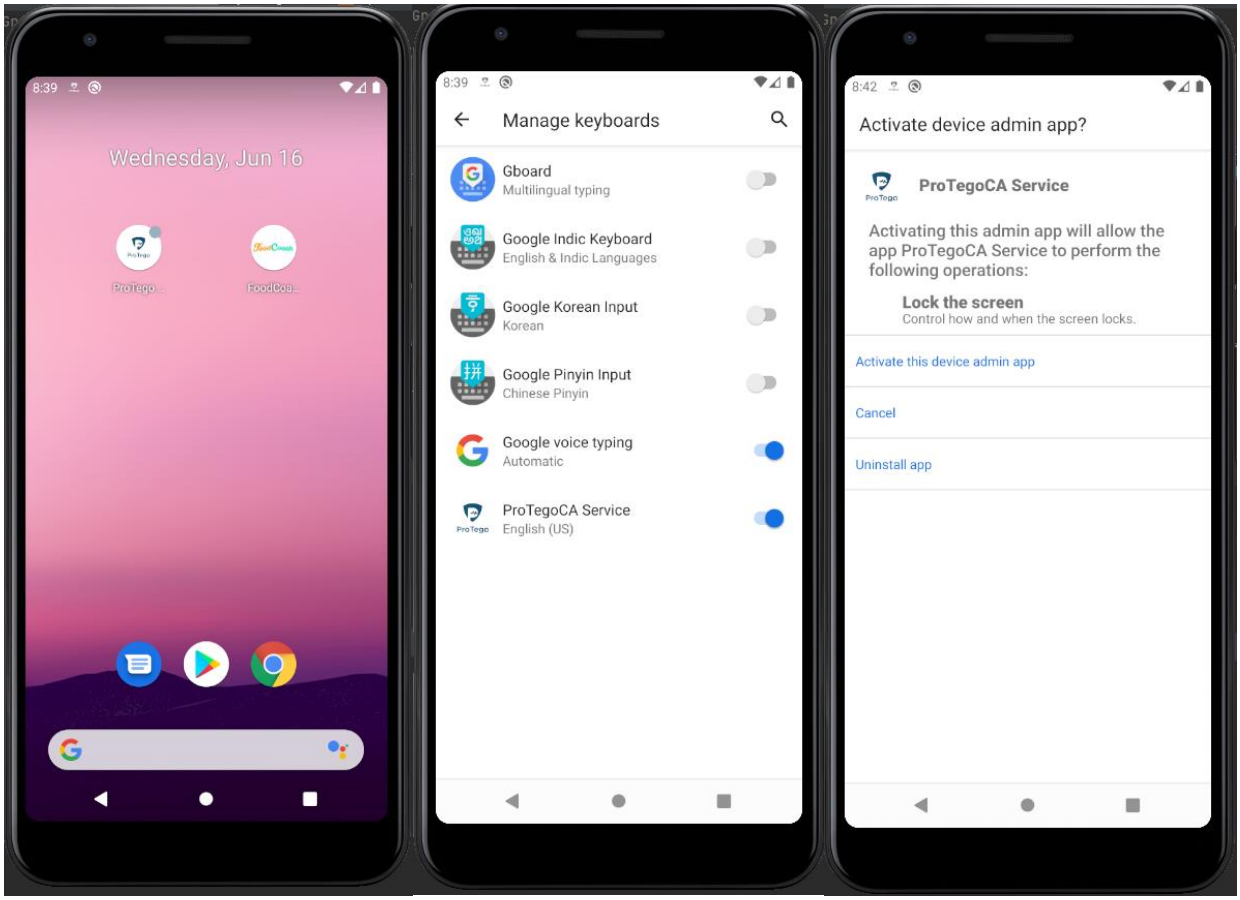

ID: S14.2

Name: Mobile device already configured

Priority: Medium

Evaluation: ✕ Failed

Description: There is no dedicated feature that allows to check whether the configuration steps have been already performed or not, other than going through the configuration again.

AT15: Store initial medical data

ID: S15.1

Name: Data operator stores initial medical data

Priority: Low

Evaluation: ✓ Passed

Description: The initial data can be provided to the application, which can seed it into the Data Gateway on startup. The screenshots below have been cropped for the sake of brevity.

```
data-gateway-init_1   No observation present in the parquet files
data-gateway-init_1   Getting the IAM Token for Elisa
data-gateway-init_1   Got the token
data-gateway-init_1   Saving initial events
data-gateway-init_1   The data gateway is pre-populated now
data-gateway-init_1 exited with code 0
```

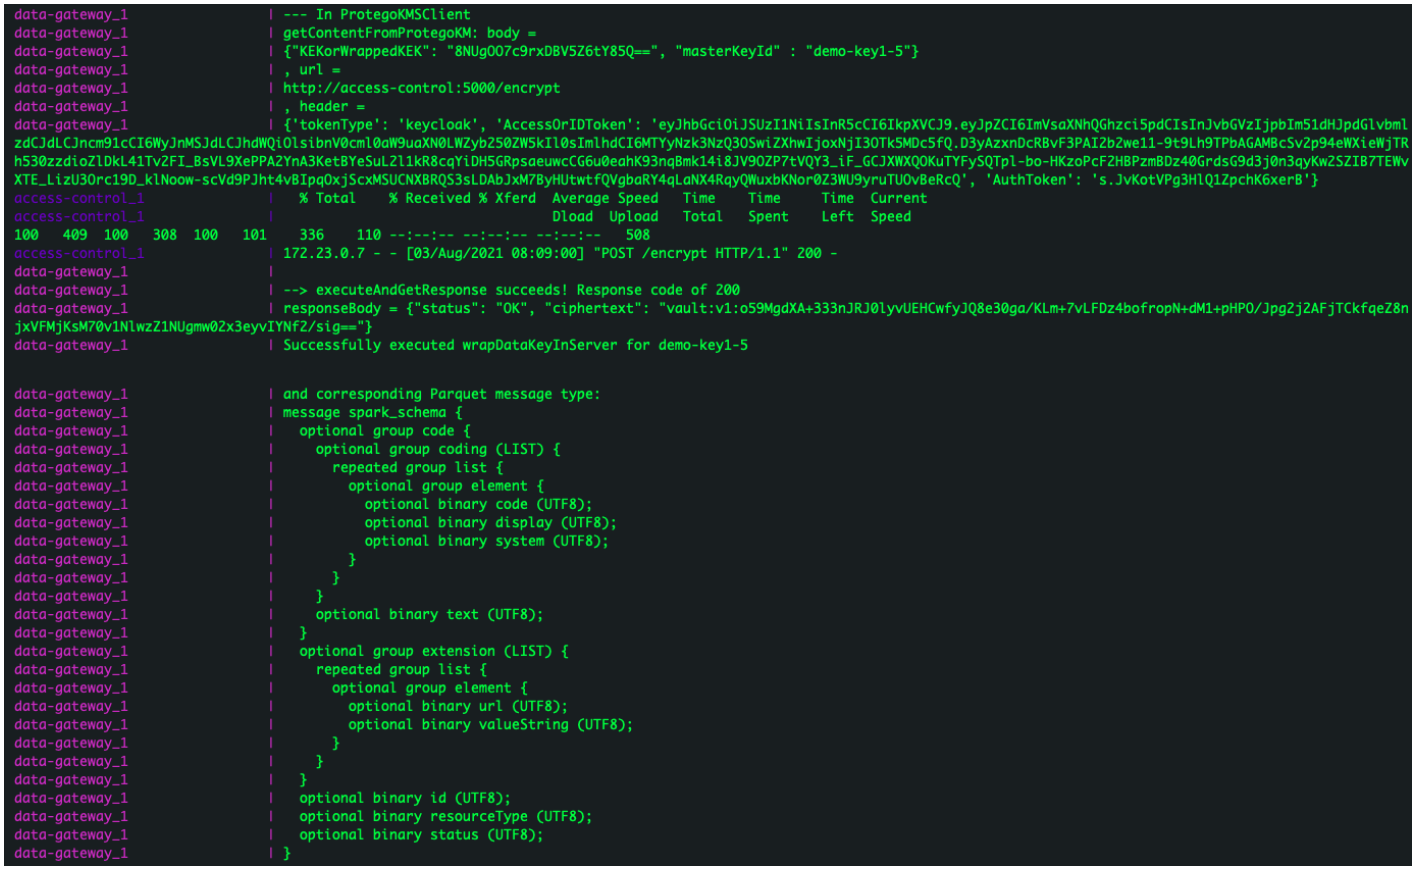

ID: S15.2

Name: Insufficient authorization

Priority: Low

Evaluation: ✓ Passed

Description: In order to seed initial data into the Data Gateway on startup, the application needs the privileges of the Nutritionist role. Failing to authenticate prevents the application from seeding initial data into the Data Gateway on startup, as the privileges of Nutritionist are required.

```
data-gateway-init_1  Getting the IAM Token for Mario
identity-access_1    | POST /nutritionist/login 403 7.658 ms - 9
data-gateway-init_1  Failed to authenticate: Forbidden
```

ID: S15.3

Name: Application initial resources are already stored

Priority: Low

Evaluation: ✓ Passed

Description: The application (i.e., FoodCoach) currently checks if initial data has already been seeded into the data storage (i.e., Data Gateway) by performing a generic query on it and checking whether any data is present on the data storage at all. If the query returns any data at all, then the application will not try to store the initial data again. In a more mature development of the toolkit, the application developers and the data storage provider could agree on an API that specifically marks this set of data as “initial data”, rather than storing it as generic data.

```
data-gateway-init_1            | Overwrite?  "no"
data-gateway-init_1            | Data are already present
dev_data-gateway-init_1 exited with code 0
```

```
await backOff(async function () {
    const observations = await issueQuery(query);
    console.log("Overwrite? ", process.env.OVERWRITE_OBSERVATIONS);
    if (observations.length > 0 && process.env.OVERWRITE_OBSERVATIONS != 'yes') {
        console.log("Data are already present");
        return;
    }
```

AT16: Register mobile device

ID: S16.1
Name: Network operator registers a new mobile device
Priority: High
Evaluation: ✓ Passed
Description: The tester logged in to the Network Slicing controller and registered a mobile device by providing its MAC Address and a name.

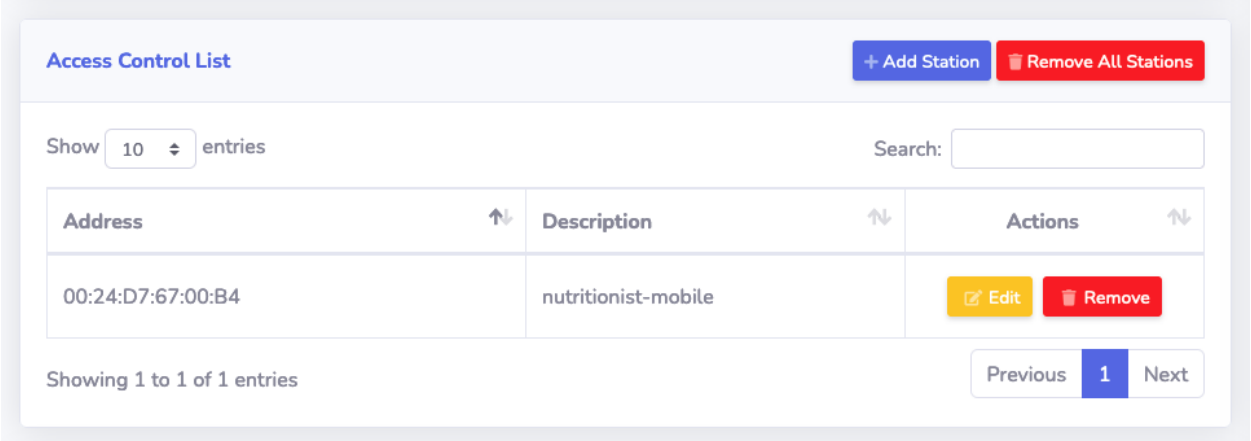

ID: S16.2
Name: Insufficient authorization
Priority: Medium
Evaluation: ✓ Passed
Description: The testers tried to log in to the Network Slicing controller with credentials that were not associated with a Network operator, and the operation failed.

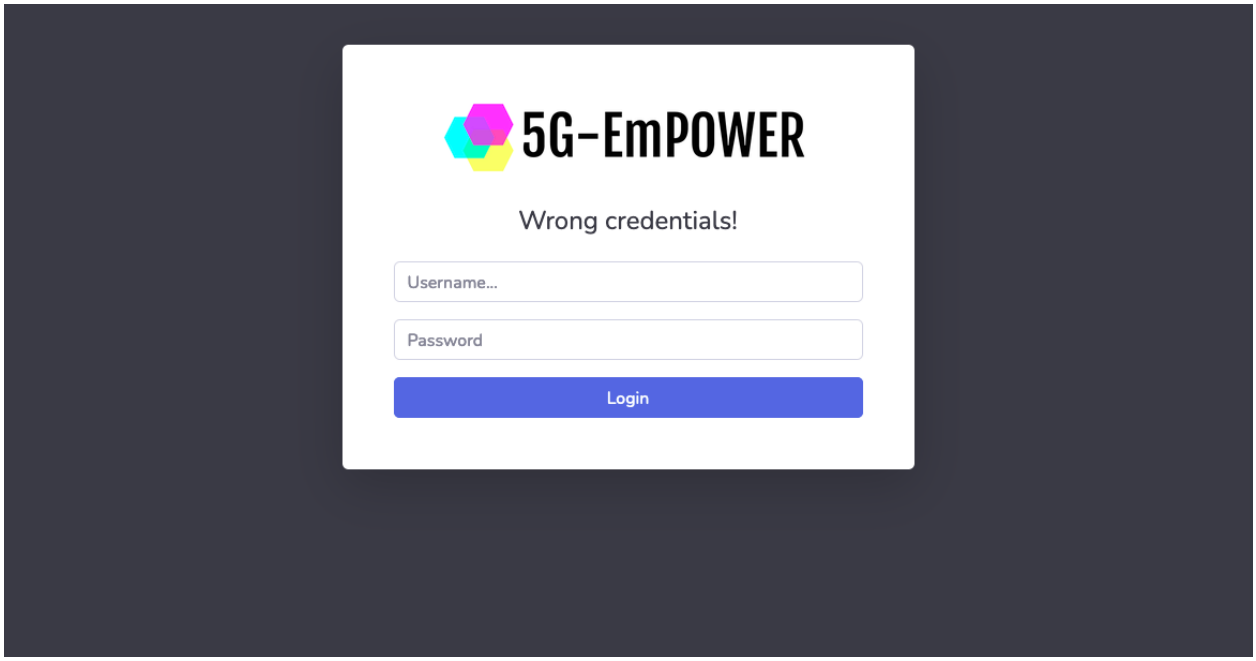

ID: S16.3
Name: Mobile device is already registered
Priority: Medium
Evaluation: ✓ Passed
Description: The testers tried to register a device with an address that had already been used for another registered device, and the operation failed, as expected.

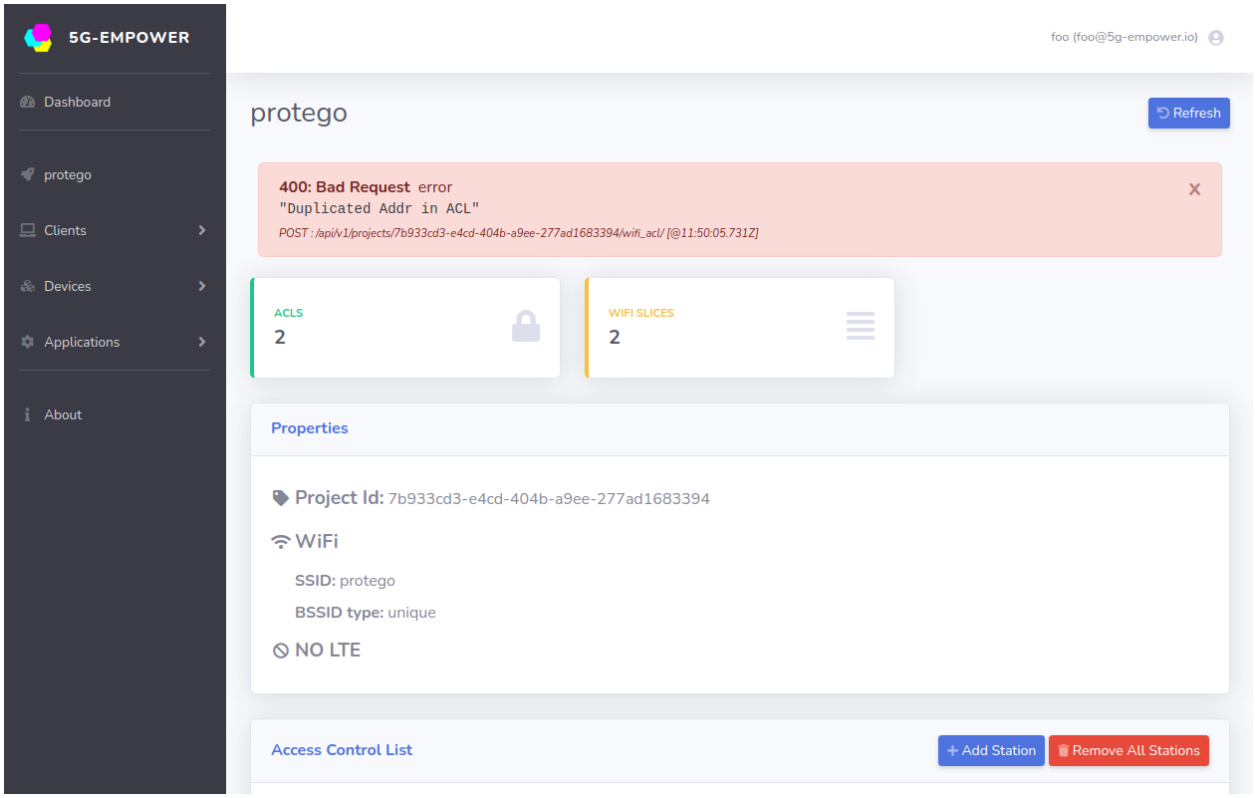

AT17: Store medical data

ID: S17.1

Name: Nutritionist stores patients' medical data

Priority: High

Evaluation: ✓ Passed

Description: The testers, acting as the Nutritionist, got the authorization token and stored the patient's records in ProTego.

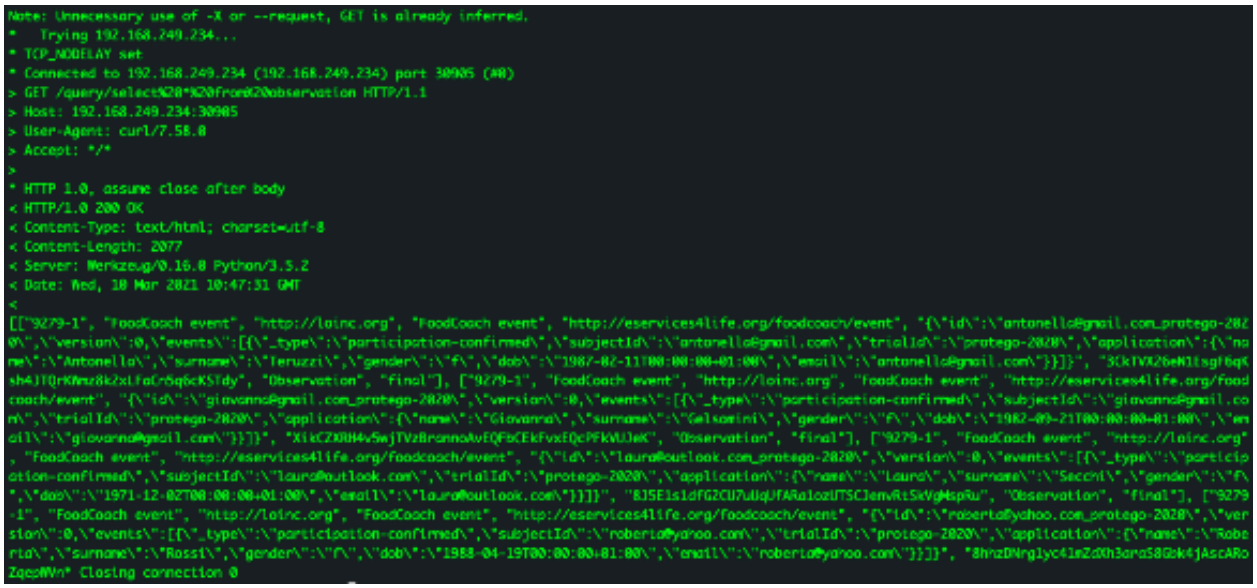

ID: S17.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The tester, acting as an external user, tried to store the patient's records in ProTego and the request was rejected.

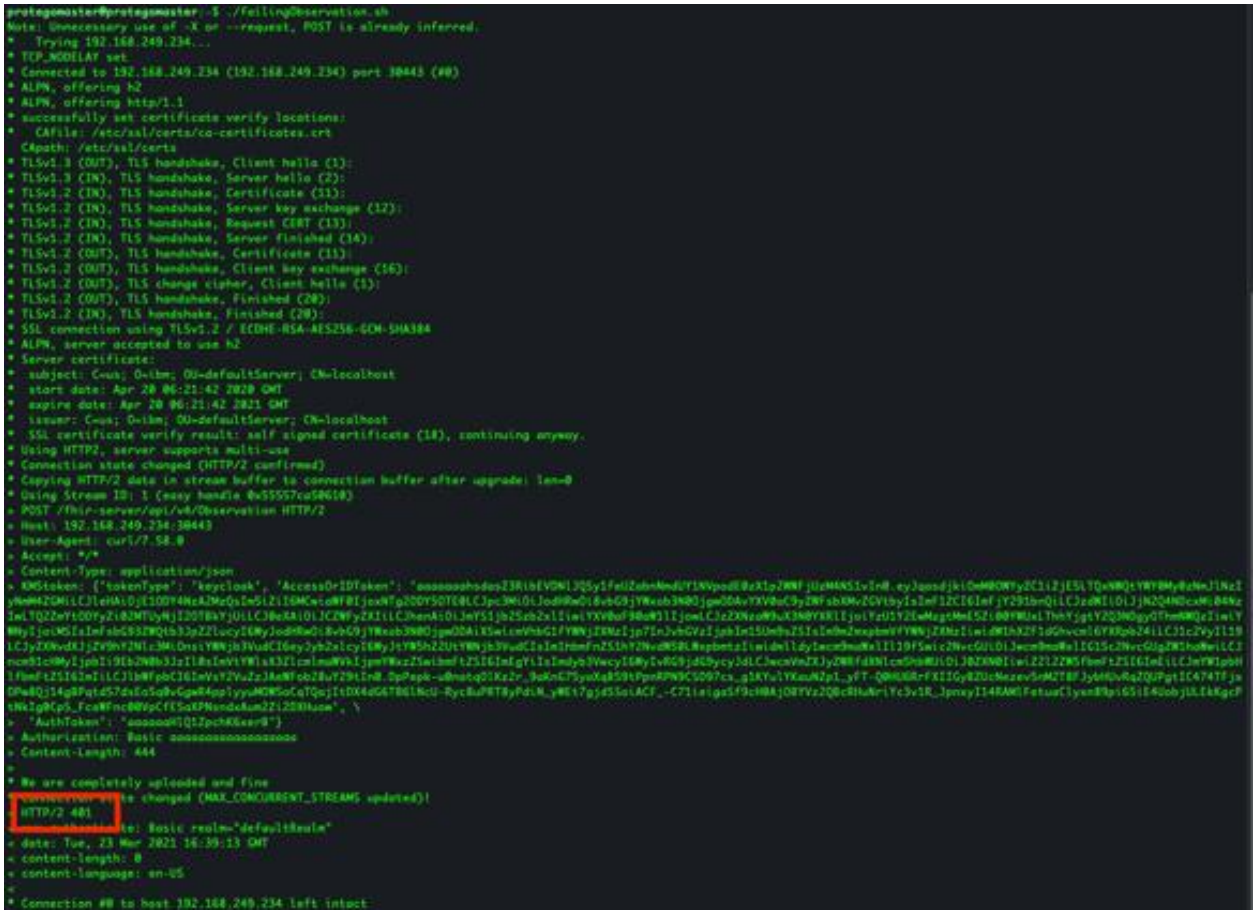

ID: S17.3

Name: Authorization token is expired

Priority: Medium

Evaluation: ✓ Passed

Description: The testers attempted to store data by using an expired token, and the request was rejected as expected.



ID: S17.4

Name: Authorization token is ill-formed

Priority: Medium

Evaluation: ✓ Passed

Description: The testers, impersonating the nutritionist, forged an ill-formed token. When they tried to store a new record, the request was rejected.

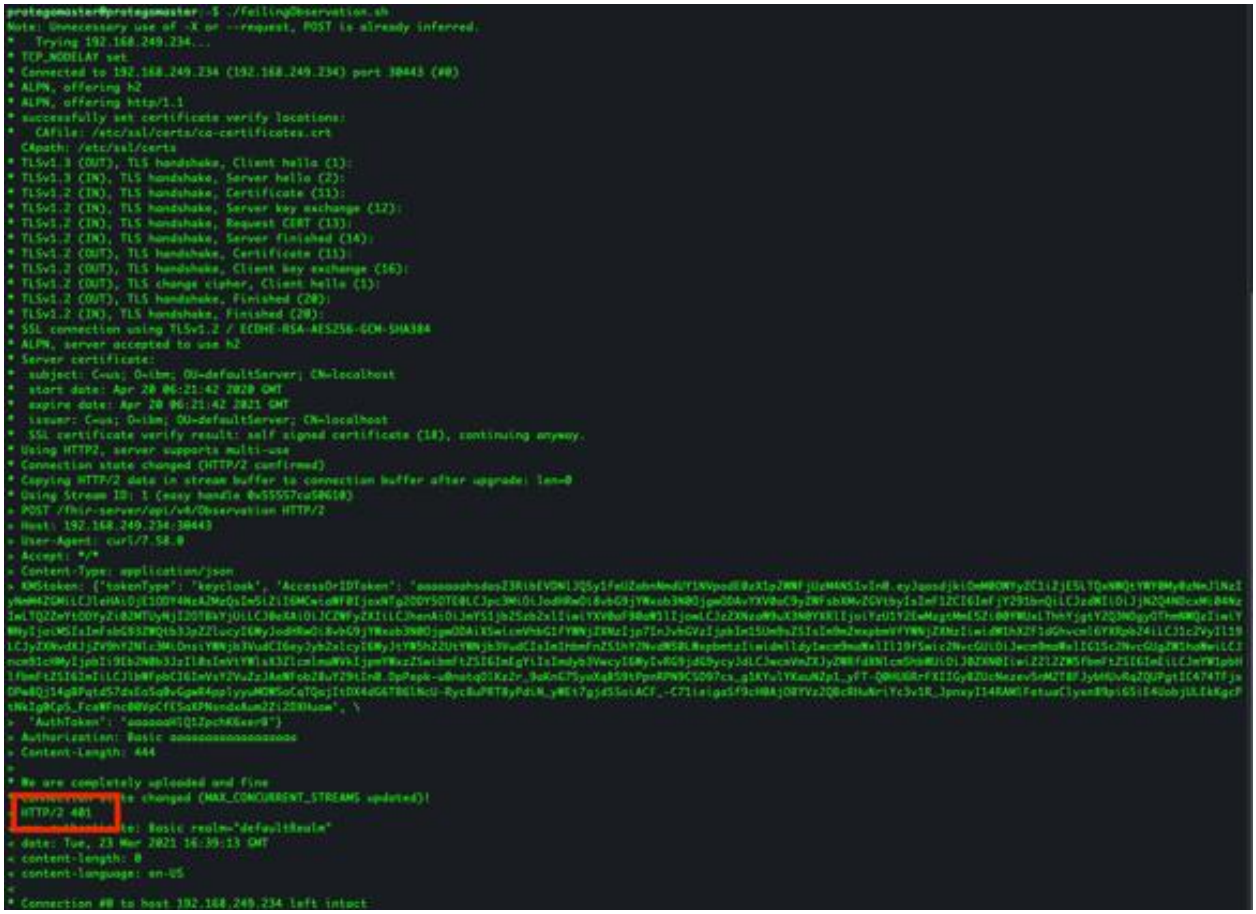

AT18: Retrieve medical data

ID: S18.1
Name: Nutritionist queries for patients' data
Priority: High
Evaluation: ✓ Passed
Description: The testers, acting as the Nutritionist, got the authorization token and retrieved the patient's records in ProTego.


```
➜ Desktop ./get-request.sh
Note: Unnecessary use of -X or --request, GET is already inferred.
*   Trying 127.0.0.1...
* TCP_NODELAY set
* Connected to localhost (127.0.0.1) port 9005 (#0)
> GET /query/select%20*%20from%20observation HTTP/1.1
> Host: localhost:9005
> User-Agent: curl/7.54.0
> Accept: */*
> KMStoken: {'tokenType': 'keycloak', 'AccessOrIDToken': 'eyJhbGciOiJSUzI1NiIsInR5cCIgOiAiSldUIiwia2lkIiA6ICJXNWtWZ3RibEVDNlJQSy1feUZabnNmdUY1NVpodE0zX1p2WWFjUzM4NS1vIn0.eyJqdGkiOiJhNmM0OWYyZC1iZjE5LTQxNWQtYWY0My0zNmJlNzIyNmM4ZGMiLCJleHAiOjE1ODY4NzAZMzQsIm5iZiI6MCwiaWF0IjoxNTg2ODY5OTE0LCJpc3MiOiJodHRwOi8vbG9jYWxob3N0OjgwODAvYXV0aC9yZWFsbXMvZGVtbyIsImF1ZCI6ImFjY291bnQiLCJzdWIiOiJjN2Q4NDcxMi04NzIwLTQ2ZmYtODYyZi02MTUyNjI2OTBkYjUiLCJ0eXAiOiJCZWFyZXIiLCJhenAiOiJCZWFyZXIiLCJzZXNzaW9uX3N0YXRlIjoiYzU1Y2EwMzgtMmE5Zi00YWUxLThhYjgtY2Q3NDgyOThmNWQzIiwiYWNyIjoiMSIsImFsbG93ZWQtb3JpZ2lucyI6WyJodHRwOi8vbG9jYWxob3N0OjgwODAiXSwicmVhbG1fYWNjZXNzIjp7InJvbGVzIjpbIm1Um9sZSIsIm9mZmxpbmVfYWNjZXNzIiwidW1hX2F1dGhvcml6YXRpb24iLCJ1c2VyIl19LCJyZXNvdXJjZV9hY2Nlc3MiOnsiYWNjb3VudCI6eyJyb2xlcyI6WyJtYW5hZ2UtYWNjb3VudCIsIm1hbmFnZS1hY2NvdW50LWxpbmtzIiwidmlldy1wcm9maWxlIl19fSwic2NvcGUiOiJwcm9maWxlIG15c2NvcGUgZW1haWwiLCJncm91cHMiOlsiZGVtbyJdLCJlbWFpbF92ZXJpZmllZCI6ZmFsc2UsImFtZSI6IEdyYiIsImdyb3VwcyI6WyIvRG9jdG9ycyJdLCJwcmVmZXJyZWRfdXNlcmShbWUiOiJ0ZXN0IiwiZ2l2ZW5fbmFtZSI6IkEiLCJmYW1pbHlfbmFtZSI6ImIiLCJlbWFpbCI6ImVsY2VuZzJAeWFobz8uY29tIn0.DpPepk-u0natqOlKz2r_9aKnG75yaXqB59tPpnRPN9CSD97cs_g1KYulYKauNZp1_yFT-Q0HU6RrFXIIGyBZUcNezevSnM2T8FJybHUvRqZQUPgtIC474TFjxOPw8Qj14gBPqtd57dsEo5q0vGgwR4pplyyuMOW5oCqTQajItDX4dG6TB6lNcU-Ryc8uPRT8yPdiN_yWEt7gjdSSoiACF_-C71ieigaSf9cH0AjOBYVz2QBcRHuNriYc3v1R_JpnxyI14RAMlFetuaClyxn89pi65iE4UobjULEkKgcPtNkIg0CpS_FceWFnc00VpCfE5aXPNsndxAum2ZiZDXHuow', 'AuthToken': 's.JvKotVPg3HlQ1ZpchK6xerB'}
> Authorization: Basic ZmhpcnVzZXI6ZmhpcnVzZXI=
> Content-Type: application/json
>
* HTTP 1.0, assume close after body
< HTTP/1.0 200 OK
< Content-Type: text/html; charset=utf-8
< Content-Length: 2641
< Server: Werkzeug/0.16.0 Python/3.5.2
< Date: Mon, 02 Aug 2021 08:05:50 GMT
<
* Closing connection 0
[[[[["9279-1", "FoodCoach event", "http://loinc.org"]], "FoodCoach event"], [["http://eservices4life.org/foodcoach/event", "{\"id\":\"antonella@gmail.com_protego-2020\",\"version\":0
,\"events\":[{\"_type\":\"participation-confirmed\",\"subjectId\":\"antonella@gmail.com\",\"trialId\":\"protego-2020\",\"application\":{\"name\":\"Antonella\",\"surname\":\"Teruzzi\"
,\"gender\":\"f\",\"dob\":\"1987-02-11T00:00:00+01:00\",\"email\":\"antonella@gmail.com\"}}]}"]], "3CkTVX26eN1EsgF6qKsh4JTQrKWmz8k2xLFaCn5q6cKSTdy", "Observation", "final"], [[[["927
9-1", "FoodCoach event", "http://loinc.org"]], "FoodCoach event"], [["http://eservices4life.org/foodcoach/event", "{\"id\":\"giovanna@gmail.com_protego-2020\",\"version\":0,\"events\
":[{\"_type\":\"participation-confirmed\",\"subjectId\":\"giovanna@gmail.com\",\"trialId\":\"protego-2020\",\"application\":{\"name\":\"Giovanna\",\"surname\":\"Gelsomini\",\"gender\
":\"f\",\"dob\":\"1982-09-21T00:00:00+01:00\",\"email\":\"giovanna@gmail.com\"}}]}"]], "XikCZXRH4v5wjTVzBrannoAvEQFbCEkFvxEQcPFkVUJeK", "Observation", "final"], [[[["9279-1", "FoodCo
ach event", "http://loinc.org"]], "FoodCoach event"], [["http://eservices4life.org/foodcoach/event", "{\"id\":\"laura@outlook.com_protego-2020\",\"version\":0,\"events\":[{\"_type\":
\"participation-confirmed\",\"subjectId\":\"laura@outlook.com\",\"trialId\":\"protego-2020\",\"application\":{\"name\":\"Laura\",\"surname\":\"Secchi\",\"gender\":\"f\",\"dob\":\"197
1-12-02T00:00:00+01:00\",\"email\":\"laura@outlook.com\"}}]}"]], "8J5E1s1dfG2CU7uUqUfARa1ozUTSCJenvRtSkVgMspRu", "Observation", "final"], [[[["9279-1", "FoodCoach event", "http://loi
nc.org"]], "FoodCoach event"], [["http://eservices4life.org/foodcoach/event", "{\"id\":\"roberta@yahoo.com_protego-2020\",\"version\":0,\"events\":[{\"_type\":\"participation-confirm
ed\",\"subjectId\":\"roberta@yahoo.com\",\"trialId\":\"protego-2020\",\"application\":{\"name\":\"Roberta\",\"surname\":\"Rossi\",\"gender\":\"f\",\"dob\":\"1988-04-19T00:00:00+01:00
\",\"email\":\"roberta@yahoo.com\"}}]}"]], "8hhzDNrg1yc41mZdXh3araS8Gbk4jAscARoZqepWVnqm", "Observation", "final"], [[[["9279-1", "FoodCoach event", "http://loinc.org"]], "FoodCoach
event"], [["http://eservices4life.org/foodcoach/event", "{\"id\":\"roberta@yahoo.com_protego-2020\",\"version\":0,\"events\":[{\"_type\":\"participation-confirmed\",\"subjectId\":\"r
oberta@yahoo.com\",\"trialId\":\"protego-2020\",\"application\":{\"name\":\"Roberta\",\"surname\":\"Rossi\",\"gender\":\"f\",\"dob\":\"1988-04-19T00:00:00+01:00\",\"email\":\"roberta
@yahoo.com\"}}]}"]], "8hhzDNrg1yc41mZdXh3araS8Gbk4jAscARoZqepWVnqm", "Observation", "final"]]%
```

ID: S18.2
Name: Insufficient authorization
Priority: Medium
Evaluation: ✓ Passed
Description: The testers, acting as an external user, created a token for an unauthorized user. Requesting the data with this token resulted in the request being rejected.

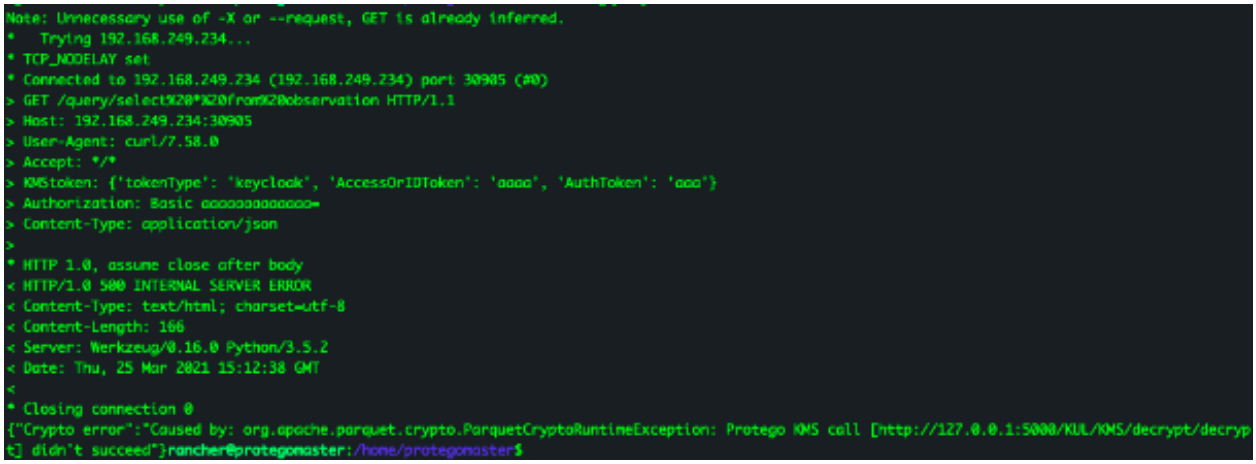

ID: S18.3
Name: Medical data have been tampered with
Priority: High
Evaluation: ✓ Passed
Description: The testers tampered with a Parquet Snappy file by changing some bytes in it. Querying the data resulted in an error, as expected.



ID: S18.4
Name: Authorization token is expired
Priority: Medium
Evaluation: ✓ Passed
Description: The testers attempted to retrieve data by using an expired token, and the request was rejected as expected.



ID: S18.5
Name: Authorization token is ill-formed
Priority: Medium
Evaluation: ✓ Passed
Description: The testers, acting as an external user, created an ill-formed token. Requesting the data with this token resulted in the request being rejected.

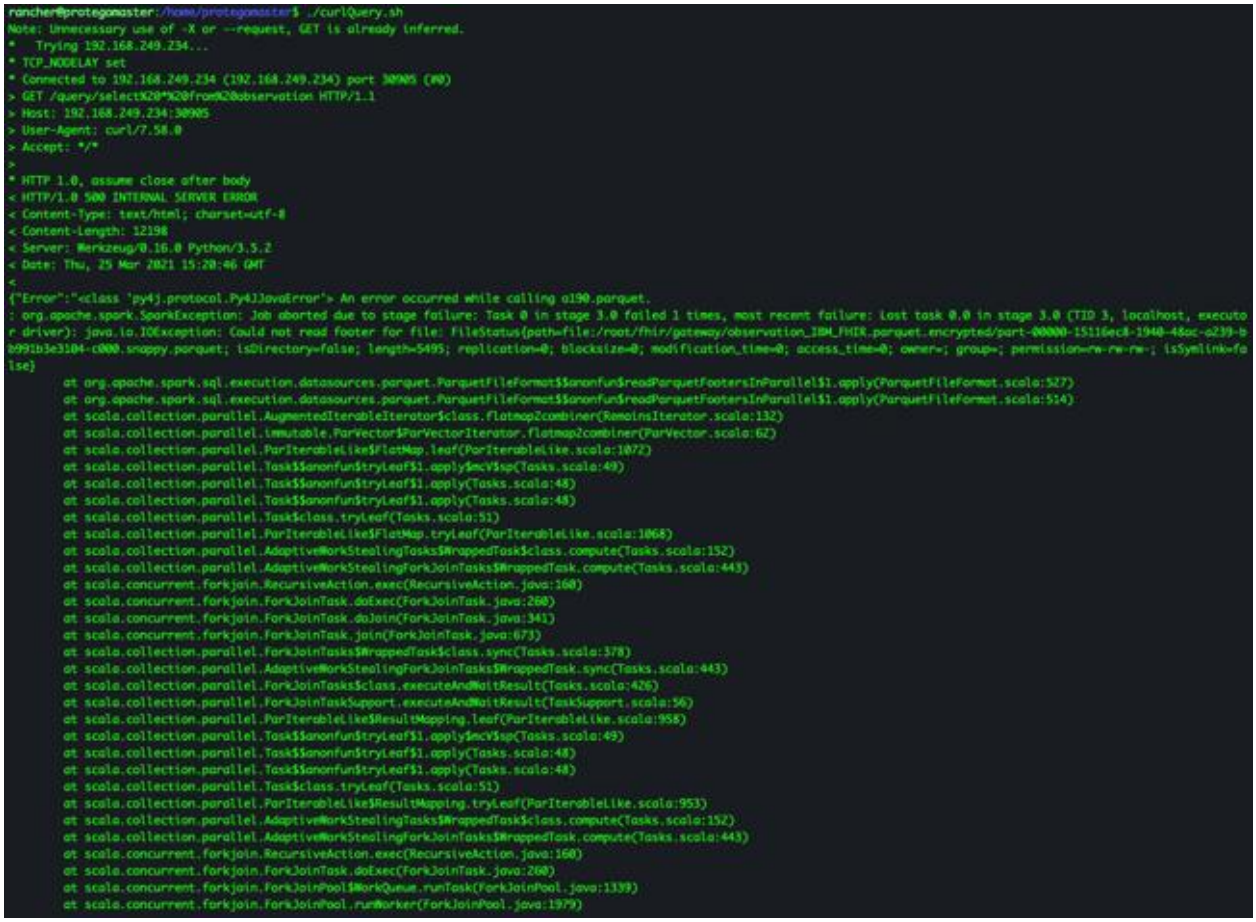

AT19: Assign IoT device to application user

ID: S19.1
Name: Patient is assigned a new IoT device
Priority: High
Evaluation: ✓ Passed
Description: The testers verified that, when valid credentials are entered by a user, the IoT device is then assigned to that user.



ID: S19.2
Name: Assign IoT device to a new user
Priority: Medium
Evaluation: ✓ Passed
Description: To verify this scenario, two steps were performed. First the tester (that is, the user) asked the system to release the previous credentials. Then, as soon as the tester tried to send new data, the system prompted them to enter their credentials.



ID: S19.3
Name: Insufficient authorization
Evaluation: Outdated
Description: The action is not performed by an administrator specifically. The system prompts to enter the credentials of the user that the IoT will be linked to. If the credentials are valid, the external IAM returns a valid token that is then stored in the IoT.

ID: S19.4
Name: Authorization token is expired
Evaluation: Outdated
Description: In this step there is no token to be used, the token is created only if the credentials entered are valid.

ID: S19.5
Name: Authorization token is ill-formed
Evaluation: Outdated
Description: In this step there is no token to be used, the token is created only if the credentials entered are valid.

AT20: Log custom application event

ID: S20.1
Name: Correct application key
Priority: High
Evaluation: ✓ Passed
Description: The testers verified that the application logs the custom application events. The agent that collects the logs is registered using the TLS protocol to obtain an encryption key that is later used to ship the logs to the component analyzes them.

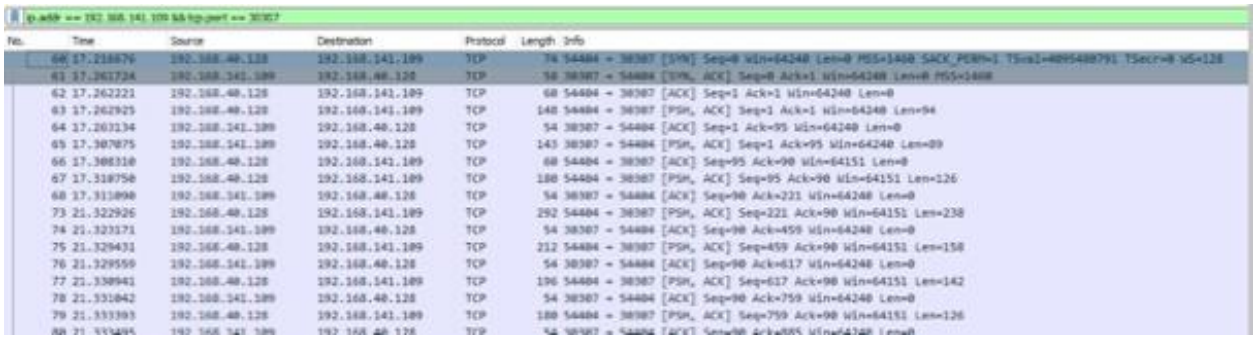

ID: S20.2
Name: Wrong application key
Priority: High
Evaluation: ✓ Passed

Description: The testers verified that any attempt of trying to ship logs without the proper encryption key is rejected.

```
2021/03/25 13:25:19 ossec-agent: WARNING: (4101): Waiting for server reply (not started). Tried: '192.168.142.109'.
2021/03/25 13:25:19 ossec-agent: WARNING: Unable to connect to any server.
2021/03/25 13:25:19 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:19 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:29 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:29 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:39 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:39 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:49 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:49 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:59 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:25:59 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:00 ossec-agent: INFO: Requesting a key from server: 192.168.142.109
2021/03/25 13:26:02 ossec-agent: ERROR: Unable to connect to 192.168.142.109:1515
2021/03/25 13:26:12 ossec-agent: WARNING: (4101): Waiting for server reply (not started). Tried: '192.168.142.109'.
2021/03/25 13:26:12 ossec-agent: WARNING: Unable to connect to any server.
2021/03/25 13:26:12 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:12 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:22 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:22 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:32 ossec-agent: INFO: Closing connection to server (192.168.142.109:30307/tcp).
2021/03/25 13:26:32 ossec-agent: INFO: Trying to connect to server (192.168.142.109:30307/tcp).
```

AT21: Send medical data

ID: S21.1

Name: IoT device sends some medical data securely

Priority: High

Evaluation: ✓ Passed

Description: Once the valid credentials have been entered and a valid token has been obtained, the IoT could send data to the Data Gateway.



ID: S21.2

Name: Authorization token is ill-formed

Priority: Medium

Evaluation: ✓ Passed

Description: If an ill-formed token is used, the Access Control returns an error, preventing the access to the data, both for read or write requests.

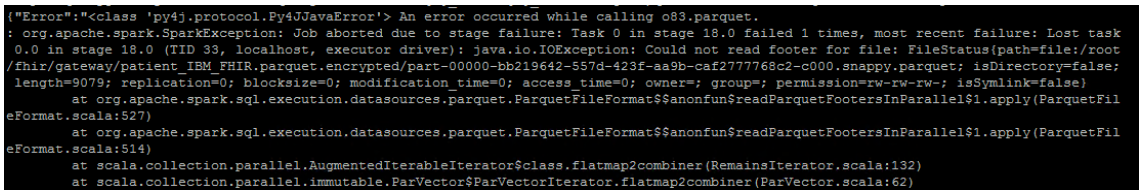

ID: S21.3

Name: Expired ID_token

Priority: Medium

Evaluation: ✓ Passed

Description: The Access Control checks if the ID_token is expired and, if so, raises an error message informing of such a situation.

ID: S21.4

Name: Protecting against Replay attacks

Priority: Medium

Evaluation: ✓ Passed

Description: The communication to reach the Data Gateway is performed with TLS, since the ELS-FHIR does only accept that kind of communication.

By having TLS in place, there is a secure connection between front-end and back-end, and any attacker will not be able to steal a valid token. Unless the attacker would be able to compromise the IoT device, but this is out of the scope of ProTego.

Taking this into account, there are two types of replay attacks one could consider:

- replay an old token: this will be detected, as the old token will have expired, and hence will be rejected by the access control framework
- replay a valid token: this is mitigated by using TLS, as described above

ID: S21.5

Name: Protecting against spoofing

Priority: Medium

Evaluation: ✓ Passed

Description: As explained in S21.4, the communication to reach the Data Gateway is performed with TLS. By using TLS, an eavesdropper cannot obtain a valid token, and therefore cannot spoof the identity of a legitimate IoT device.

AT22: Report suspicious activity

ID: S22.1
Name: Mobile agent reports suspicious activity
Priority: High
Evaluation: ✓ Passed

Description: When the level of trust of a user drops below a certain threshold, the mobile agent reports this event, both visually (on the mobile device, as shown in screenshot 1) and in the logs (in the backend, as shown in screenshot 2). Screenshot 3 shows the resulting alert that is displayed in the SIEM interface.

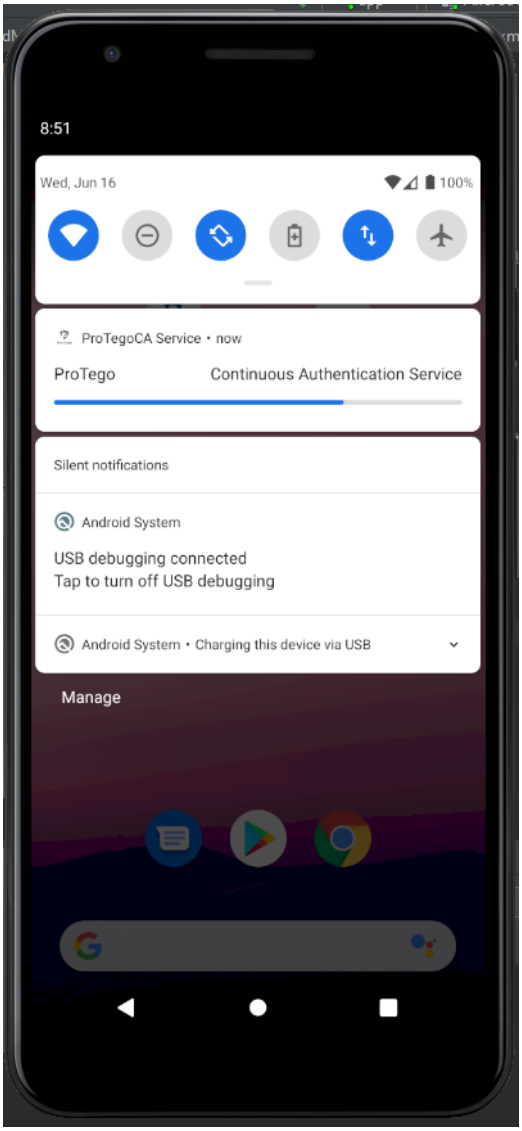


```
2021-06-16 10:49:51.330 26594-26648/eu.protego.ca.edr.android E/PMS: TRUST_DATA : 0.915
2021-06-16 10:49:51.599 26594-26648/eu.protego.ca.edr.android E/PMS: TRUST_DATA : 0.915
2021-06-16 10:49:51.867 26594-26648/eu.protego.ca.edr.android E/PMS: TRUST_DATA : 0.6
2021-06-16 10:49:55.158 26594-28137/eu.protego.ca.edr.android E/PMS: EVAL_TRUST : null
2021-06-16 10:49:55.191 26594-28138/eu.protego.ca.edr.android E/PMS: EVAL_TRUST : null
2021-06-16 10:49:55.434 26594-26648/eu.protego.ca.edr.android E/PMS: TRUST_DATA : 0.6
```

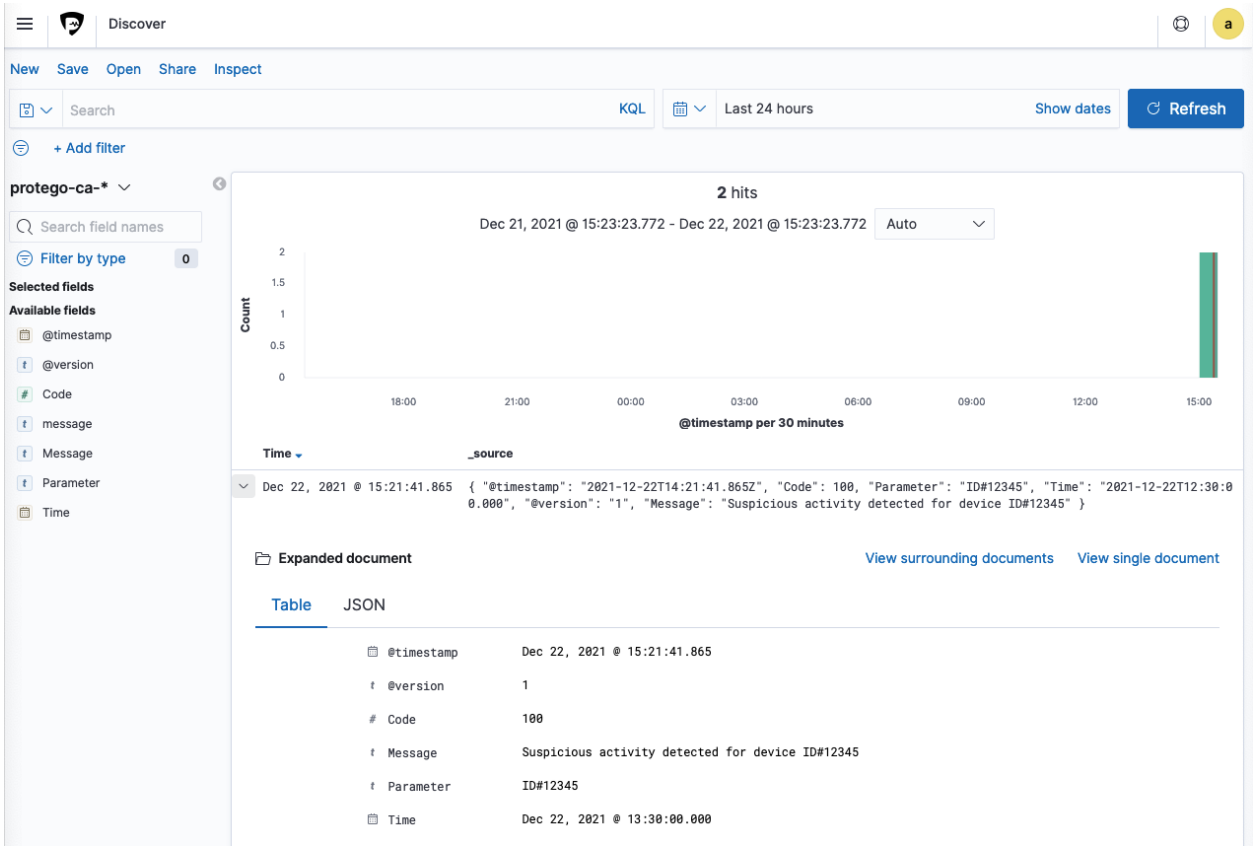
Discover
New Save Open Share Inspect
Search KQL Last 24 hours Show dates Refresh
+ Add filter
protego-ca-*
Search field names
Filter by type 0
Selected fields
Available fields
@timestamp
@version
Code
message
Message
Parameter
Time
2 hits
Dec 21, 2021 @ 15:23:23.772 - Dec 22, 2021 @ 15:23:23.772 Auto
Count
@timestamp per 30 minutes
Time _source
Dec 22, 2021 @ 15:21:41.865 { "@timestamp": "2021-12-22T14:21:41.865Z", "Code": 100, "Parameter": "ID#12345", "Time": "2021-12-22T12:30:00.000", "@version": "1", "Message": "Suspicious activity detected for device ID#12345" }
Expanded document
View surrounding documents
View single document
Table JSON
@timestamp Dec 22, 2021 @ 15:21:41.865
@version 1
Code 100
Message Suspicious activity detected for device ID#12345
Parameter ID#12345
Time Dec 22, 2021 @ 13:30:00.000

AT23: Respond to alert

ID: S23.1
Name: Security operator responds to an alert
Priority: High
Evaluation: ✓ Passed
Description: The testers verified that it is possible to set up notifications so that the operator can be warned as soon as an alert is fired. Screenshot 1 shows an example of the notification of an alert that has been received by the operator via Microsoft Teams (though, notifications are supported via other channels, such as email, Slack, or others). After that, the operator was able to look at the details of the alert on the user interface of the SIEM, shown in screenshot 2. Moreover, screenshot 3 shows a list of the possible actions that can be automated so that they are executed as soon as a certain alert is fired. No such action has been set up in the current configuration, as these actions can cause undesired effects, like the blockage of a service.

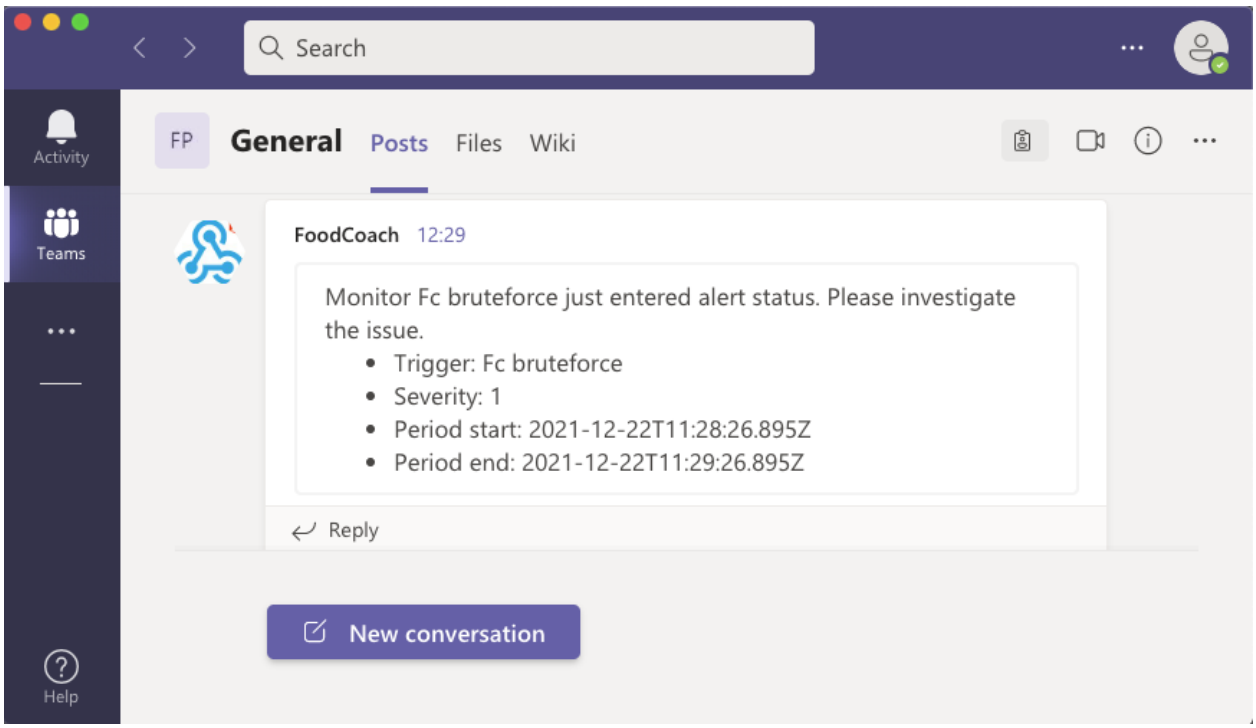

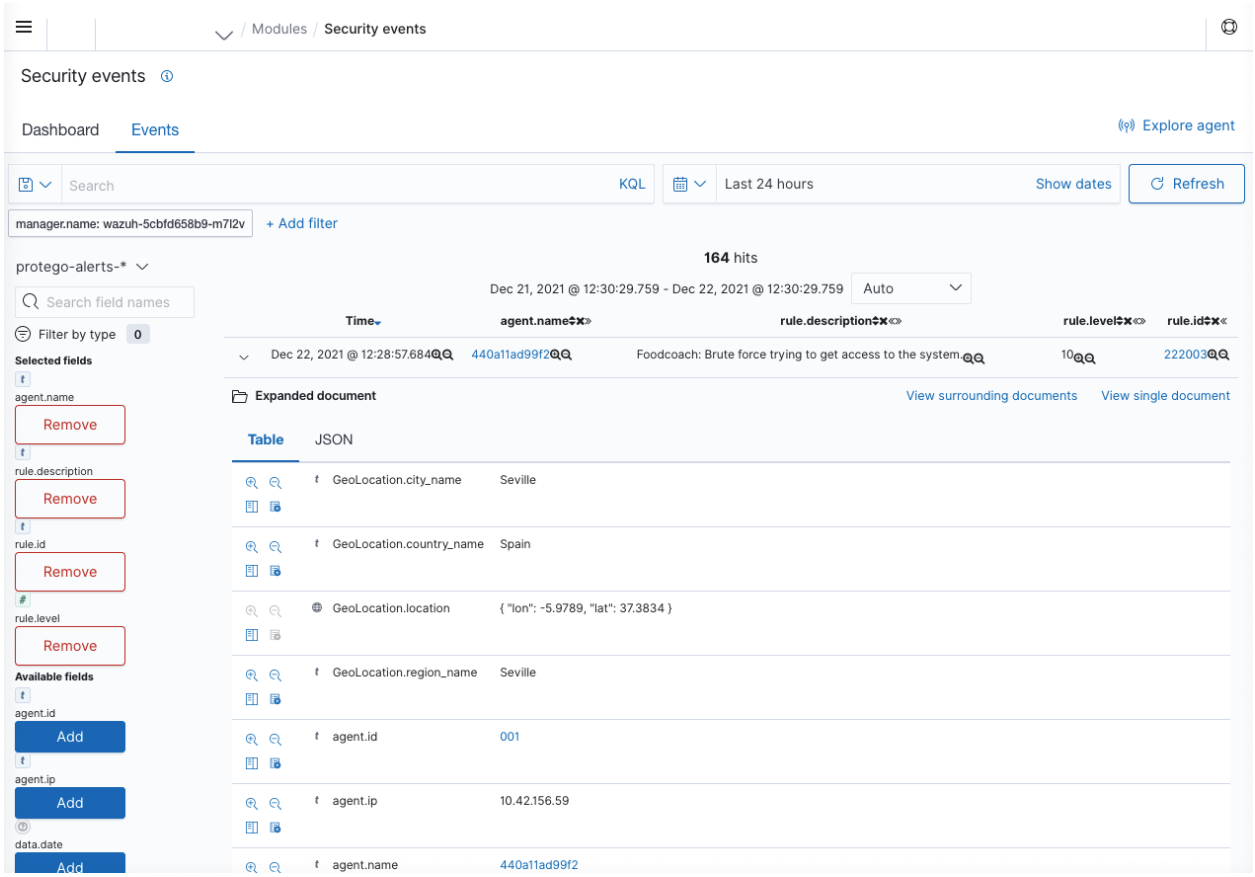
Modules / Security events
Security events
Dashboard
Events
Explore agent
Search
KQL
Last 24 hours
Show dates
Refresh
manager.name: wazuh-5cbfd658b9-m7l2v
+ Add filter
protego-alerts-*
Search field names
Filter by type 0
Selected fields
agent.name
Remove
rule.description
Remove
rule.id
Remove
rule.level
Remove
Available fields
agent.id
Add
agent.ip
Add
data.date
Add
164 hits
Dec 21, 2021 @ 12:30:29.759 - Dec 22, 2021 @ 12:30:29.759
Auto
Time
agent.name
rule.description
rule.level
rule.id
Dec 22, 2021 @ 12:28:57.684
440a11ad99f2
Foodcoach: Brute force trying to get access to the system.
10
222003
Expanded document
View surrounding documents
View single document
Table
JSON
GeoLocation.city_name
Seville
GeoLocation.country_name
Spain
GeoLocation.location
{ "lon": -5.9789, "lat": 37.3834 }
GeoLocation.region_name
Seville
agent.id
001
agent.ip
10.42.156.59
agent.name
440a11ad99f2

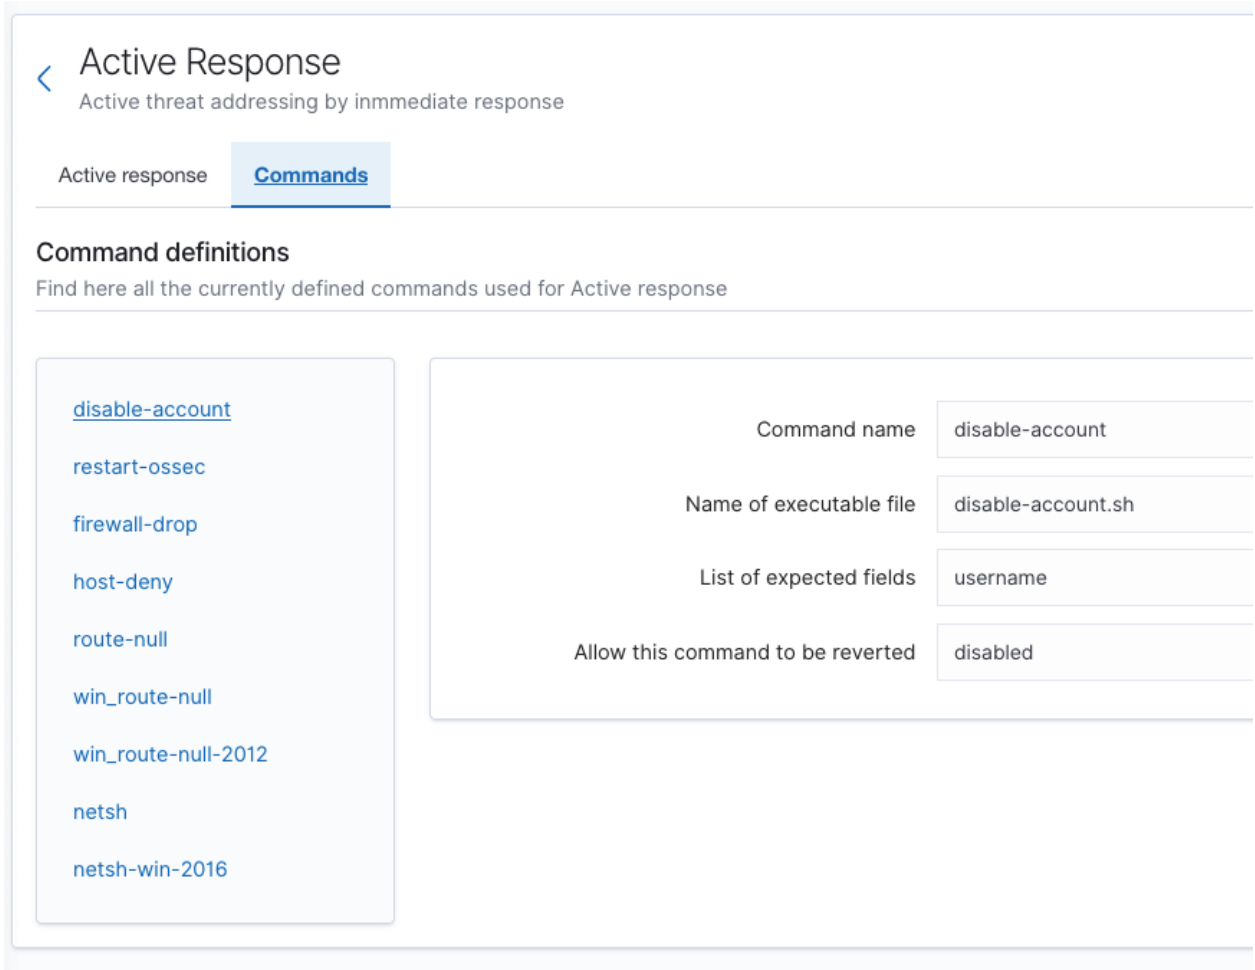
Active Response
Active threat addressing by inmmediate response
Active response
Commands
Command definitions
Find here all the currently defined commands used for Active response
disable-account
restart-ossec
firewall-drop
host-deny
route-null
win_route-null
win_route-null-2012
netsh
netsh-win-2016
Command name
disable-account
Name of executable file
disable-account.sh
List of expected fields
username
Allow this command to be reverted
disabled

ID: S23.2
Name: Insufficient authorization
Priority: Medium
Evaluation: ✓ Passed

Description: The testers verified that only authorized users can access to this functionality.

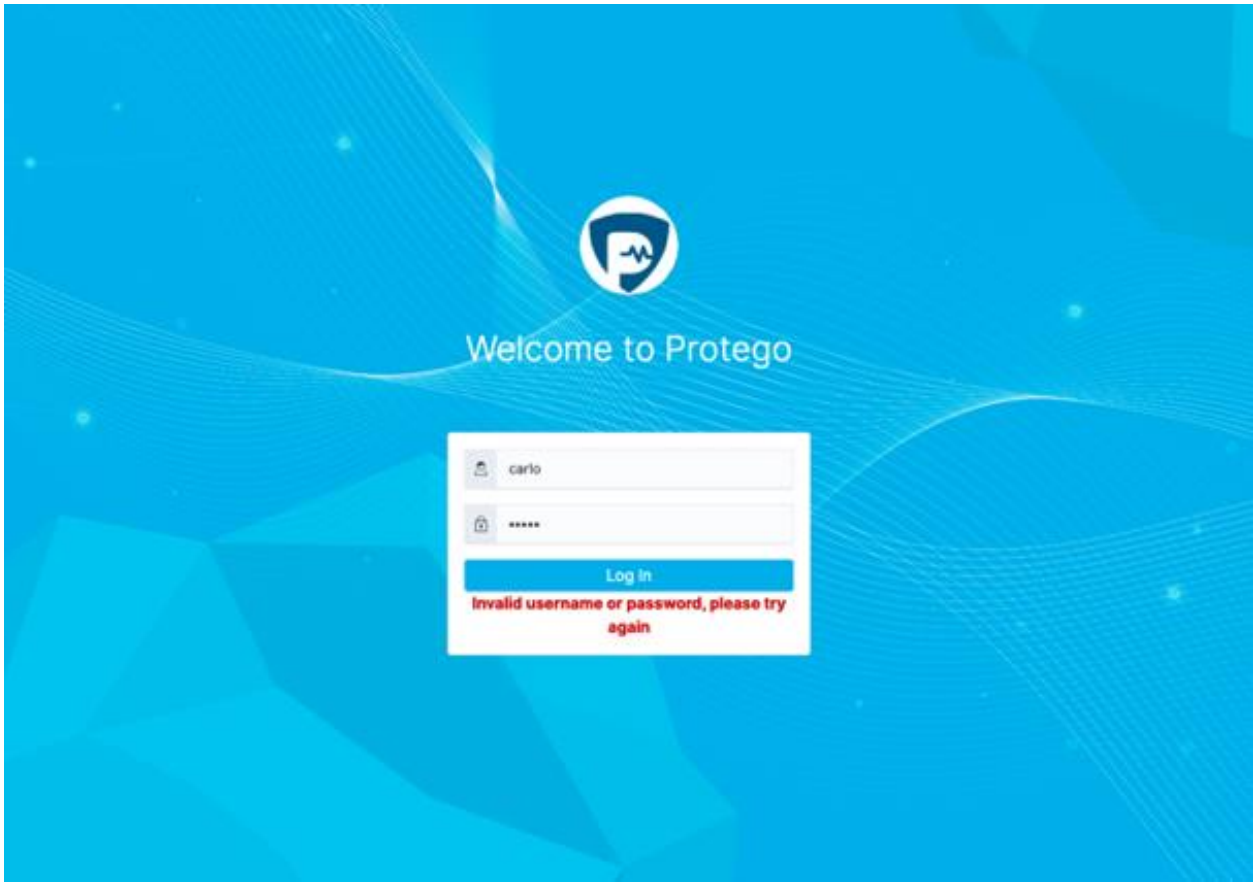

AT24: Review alerts

ID: S24.1
Name: Security operator review alerts
Priority: High
Evaluation: ✓ Passed

Description: The testers logged into the SIEM interface and requested to see an overview of the all the alerts, which is shown in the main dashboard of the interface (shown in screenshot 1). It is also possible to view a list of all the alerts that have been generated by the SIEM (shown in screenshot 2).

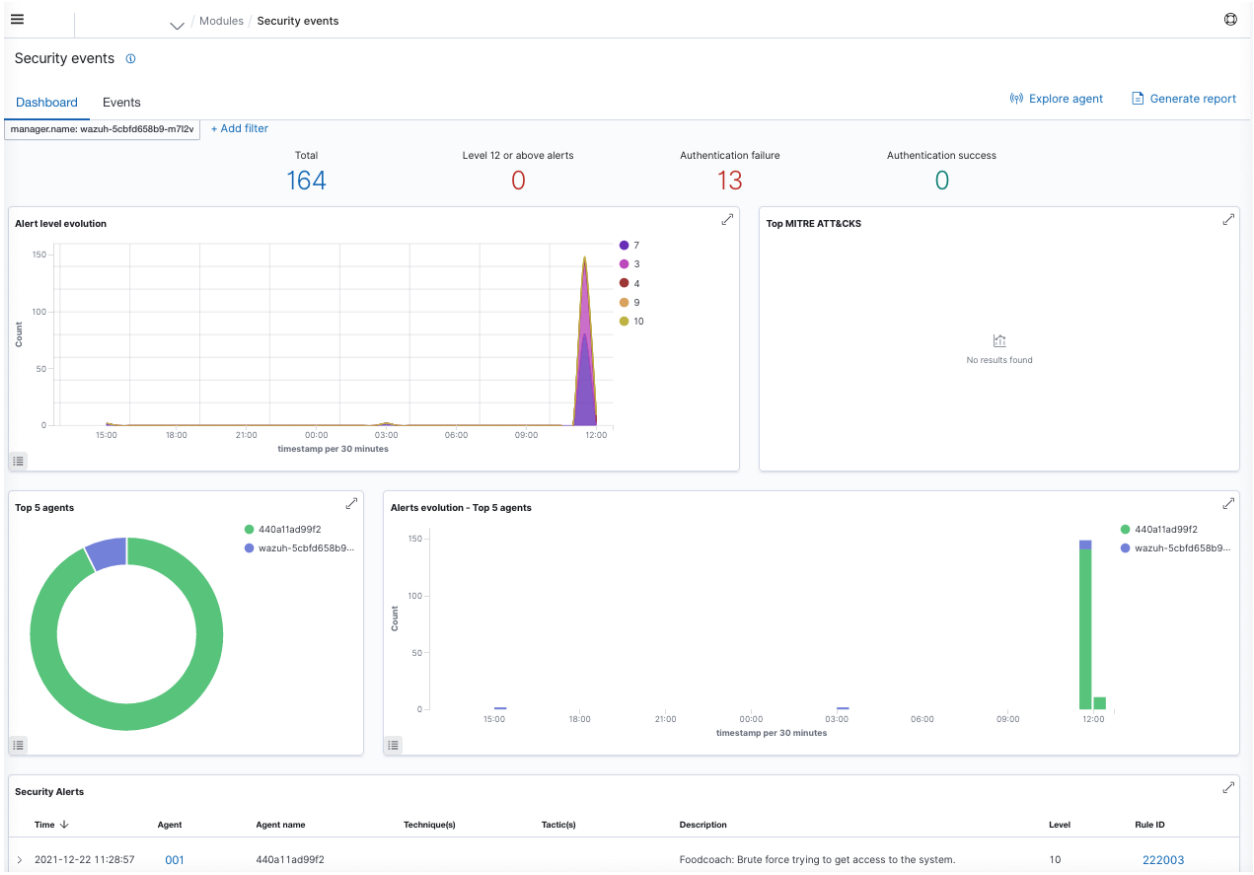

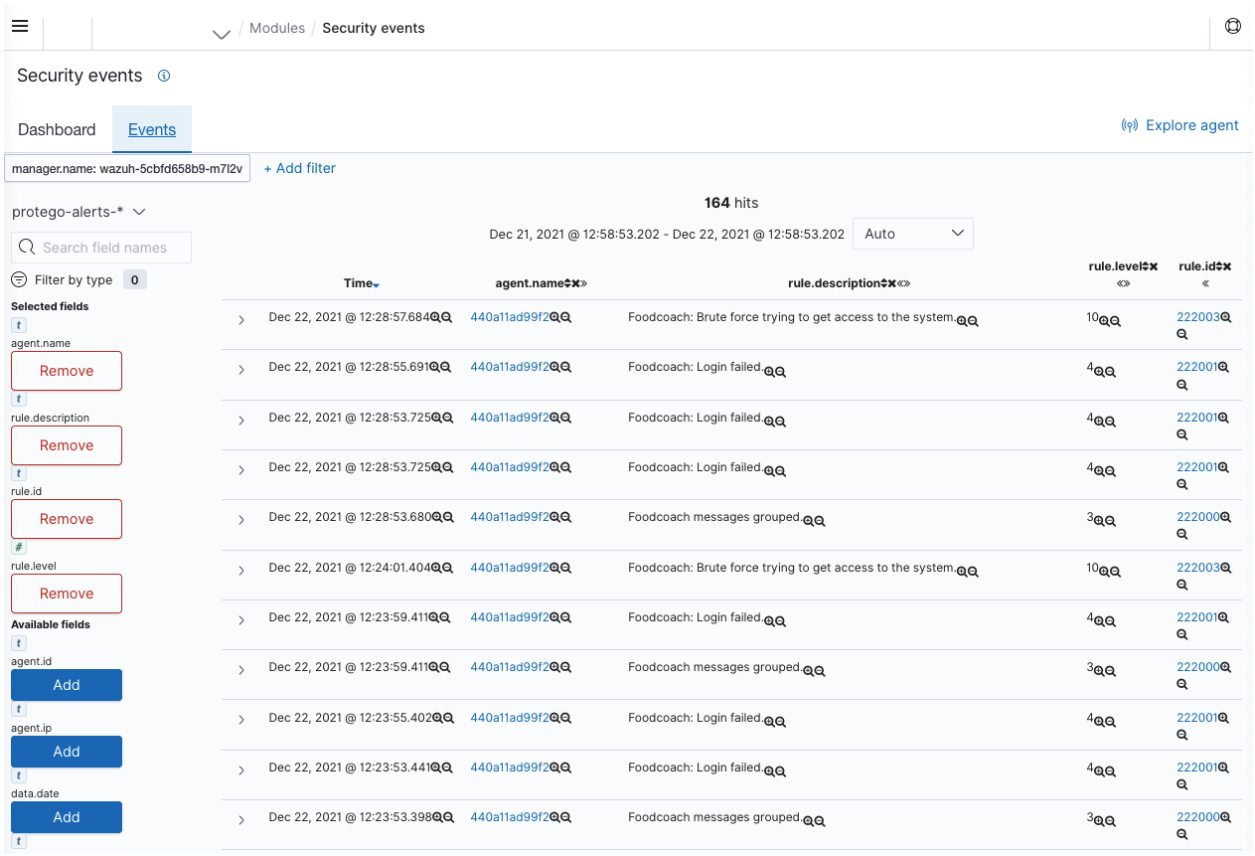
Modules / Security events
Security events
Dashboard
Events
Explore agent
manager.name: wazuh-5cbfd658b9-m7l2v
+ Add filter
protego-alerts-*
Search field names
Filter by type 0
Selected fields
agent.name
Remove
rule.description
Remove
rule.id
Remove
rule.level
Remove
Available fields
agent.id
Add
agent.ip
Add
data.date
Add
164 hits
Dec 21, 2021 @ 12:58:53.202 - Dec 22, 2021 @ 12:58:53.202
Auto
Time
agent.name
rule.description
rule.level
rule.id
Dec 22, 2021 @ 12:28:57.684 440a11ad99f2 Foodcoach: Brute force trying to get access to the system. 10 222003
Dec 22, 2021 @ 12:28:55.691 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:28:53.725 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:28:53.725 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:28:53.680 440a11ad99f2 Foodcoach messages grouped. 3 222000
Dec 22, 2021 @ 12:24:01.404 440a11ad99f2 Foodcoach: Brute force trying to get access to the system. 10 222003
Dec 22, 2021 @ 12:23:59.411 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:23:59.411 440a11ad99f2 Foodcoach messages grouped. 3 222000
Dec 22, 2021 @ 12:23:55.402 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:23:53.441 440a11ad99f2 Foodcoach: Login failed. 4 222001
Dec 22, 2021 @ 12:23:53.398 440a11ad99f2 Foodcoach messages grouped. 3 222000

ID: S24.2
Name: Insufficient authorization
Priority: Medium
Evaluation: ✓ Passed
Description: The testers verified that only authorized users can access this functionality.

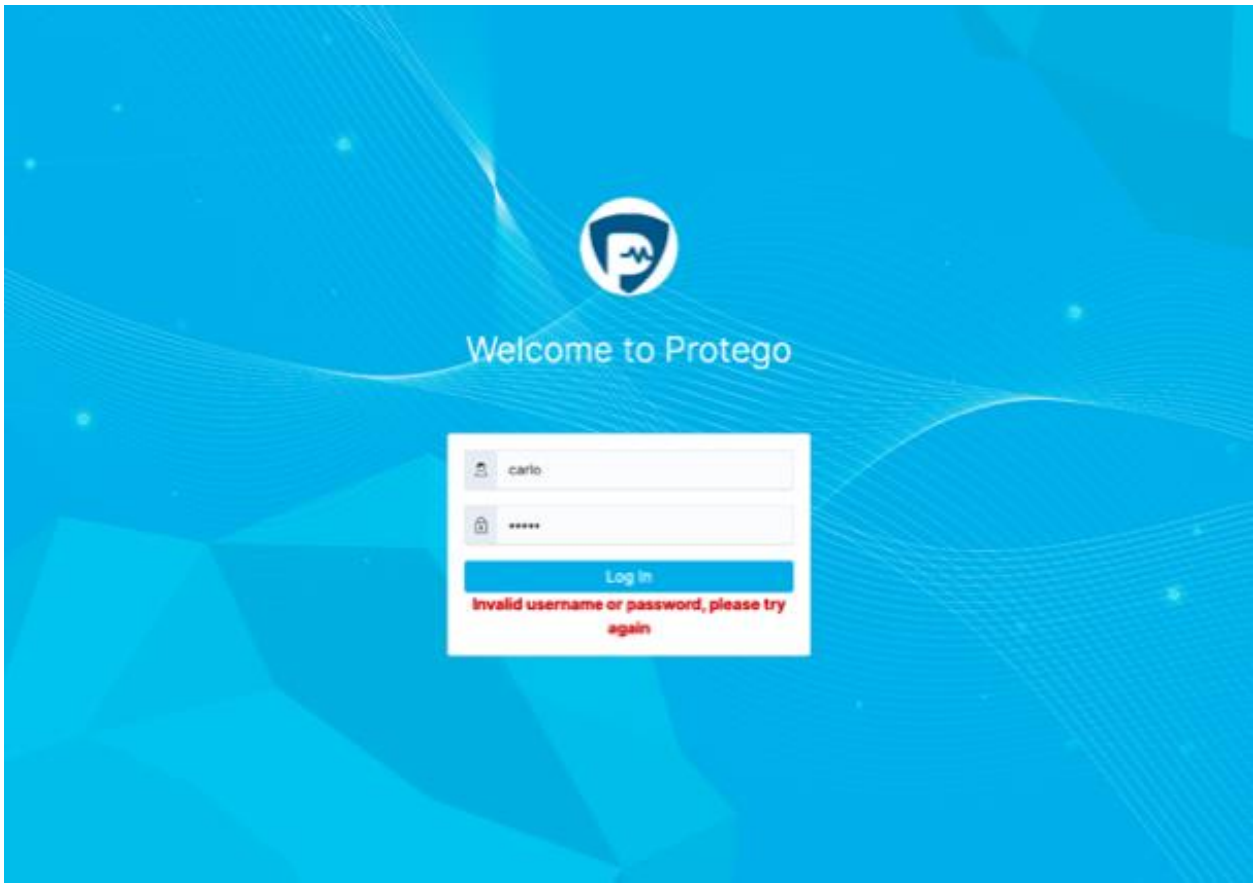

AT25: Review new risk evaluation

ID: S25.1

Name: Security operator reviews changes in the evaluation

Priority: High

Evaluation: ✓ Passed

Description: The testers verified that the SIEM can send new vulnerabilities to the SSM and that a recalculation of the risks can be performed then. Screenshot 1 shows the SSM UI before vulnerabilities are injected. As can be seen in the top right corner, the level of the highest risk is *Low*. More specifically, screenshot 2 focuses on the Trustworthiness Attributes of the FoodCoach Frontend. When new vulnerabilities are reported by the SIEM VA-Analyzer (Screenshot 3), these are communicated to the SSM, which in turn updates the Trustworthiness Attributes (Screenshot 4). Lastly, Screenshot 5 shows the outcome after a recalculation of the risks, and once again the level of the highest risk is shown in the top right corner, but this time it is *High*.

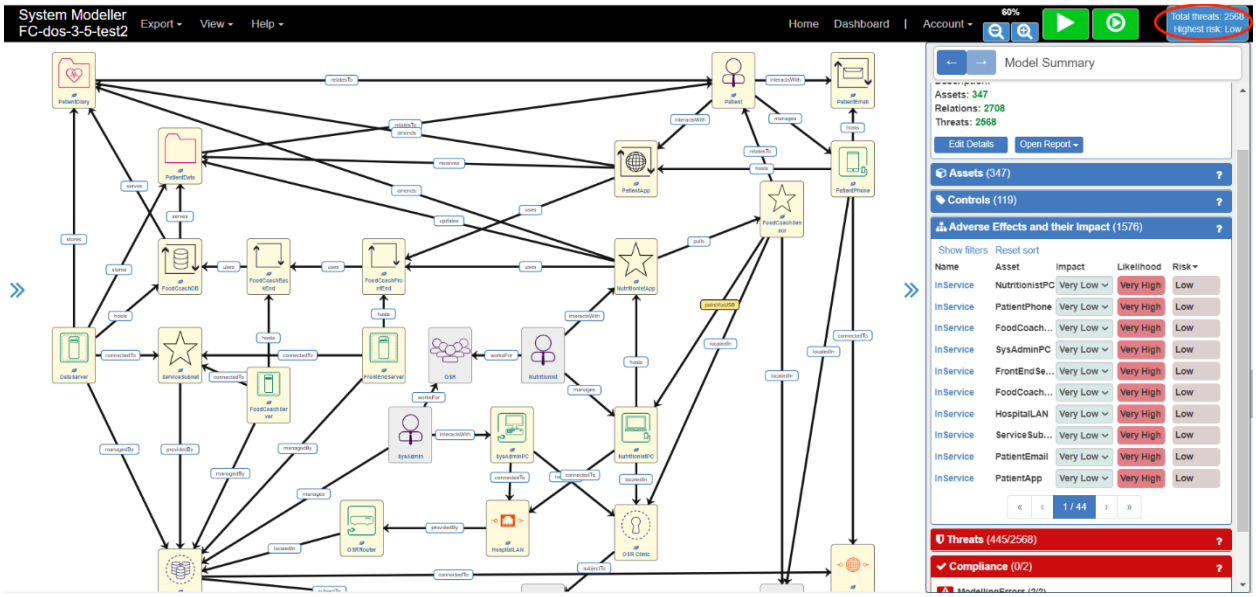

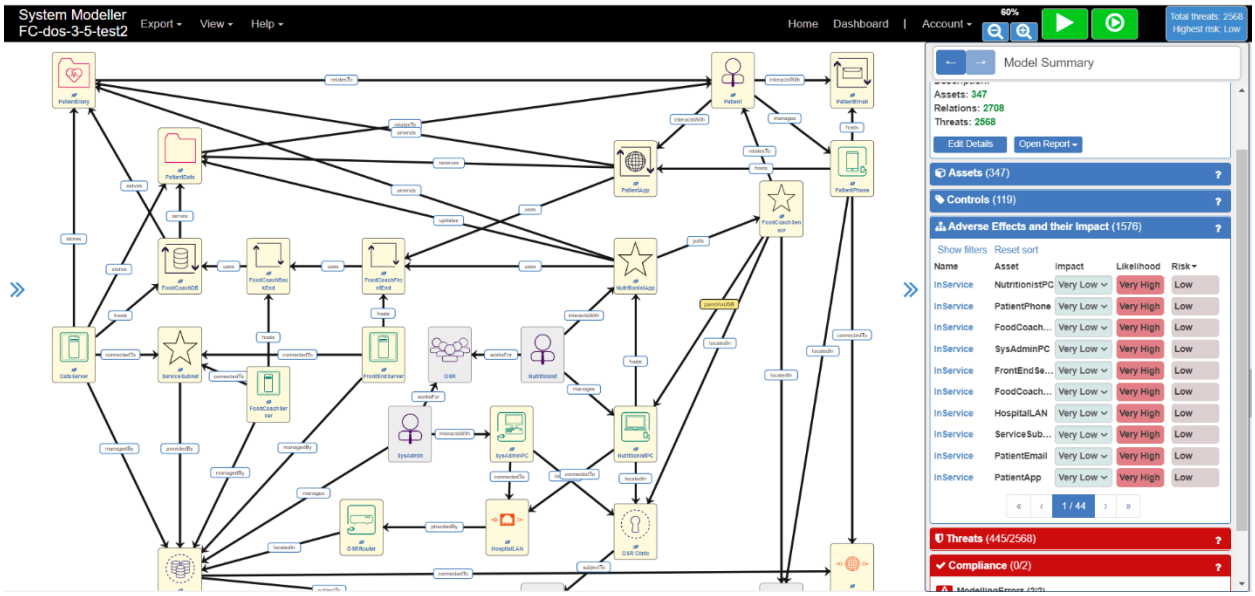


2021-03-18 11:57:23,457 - Connection status: Successful
2021-03-18 11:57:23,475 - Connected to Protego database.
2021-03-18 11:57:23,550 - Task: c22c3817-6083-4ac9-919f-4f68fa860fe1
2021-03-18 11:57:23,551 - Report: fc203df3-311d-4c85-a38d-85c624ac8704
2021-03-18 11:57:24,161 - Result: 031e1962-d436-4eac-b4a0-6b911f71839c
2021-03-18 11:57:28,519 - Result: 402cc3aa-f33a-4da6-8096-61173a11e639
2021-03-18 11:57:40,299 - Result: f864a64af-864b-4323-9c31-6d63e46bbd1d
2021-03-18 11:57:56,048 - Result: 73e380f5-d4db-4799-ab0f-a87d68c804a5
2021-03-18 11:57:56,052 - Result: 01ada523-292e-446a-9955-edae8dc1de3b
2021-03-18 11:57:56,054 - Result: d2ec0b35-0650-43d2-a198-35043eccc9f1
2021-03-18 11:58:04,964 - Result: 2d555900-ddeb-4882-b49d-6d9f9b92a28d
2021-03-18 11:58:04,968 - Result: 9a911d82-40f1-4211-9a76-82e29e6b65e0
2021-03-18 11:58:10,381 - Result: b4fce616-8201-40fd-b583-146a7a346043
2021-03-18 11:58:24,650 - Result: b509f5a6-effb-409d-885c-ac3b80c75230
2021-03-18 11:58:24,698 - Result: 414d71dc-dea2-4a96-b37f-1a5f8ce1d663
2021-03-18 11:58:24,699 - Result: d8fdc9da-aefa-4fa2-9188-9a7d1ad95fab
2021-03-18 11:58:24,700 - Result: 4ddfa124-d52e-4882-b4d7-a584919af6eb
2021-03-18 11:58:24,700 - Result: df832240-ba12-42c4-b7c8-80b43553810e
2021-03-18 11:58:24,700 - Result: 85354221-29c7-47f2-8796-47bfba741e35
2021-03-18 11:58:24,703 - Result: 629dd04d-2746-470c-ad5b-7edbc3e259ca
2021-03-18 11:58:24,703 - Result: 4f83ea2d-fffb-4662-9eab-12c78e6586f3
2021-03-18 11:58:24,705 - Result: 6a092ee5-2296-4467-b93c-72b57a708367
2021-03-18 11:58:26,530 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:58:58,451 - Vulnerability update successful!
2021-03-18 11:58:58,453 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:04,413 - Vulnerability update successful!
2021-03-18 11:59:04,415 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:10,610 - Vulnerability update successful!
2021-03-18 11:59:10,613 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:16,960 - Vulnerability update successful!
2021-03-18 11:59:16,961 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:23,736 - Vulnerability update successful!
2021-03-18 11:59:23,738 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:30,849 - Vulnerability update successful!
2021-03-18 11:59:30,851 - Updating FoodCoach Front End vulnerabilities...
2021-03-18 11:59:36,477 - Vulnerability update successful!
2021-03-18 11:59:36,479 - Recalculating 3cju27up9llbqoog52qqm2tjhlntjb8pn2tarvrbk12p8t41uumi2hd6e70n67q2cc0g92uv80acktb5kkpbqaggkhh131balktmpds risks...
2021-03-18 11:59:36,684 - Risks recalculation task created with id: 605340a8b4f841b95295725c
2021-03-18 11:59:36,685 - Checking 605340a8b4f841b95295725c status...
2021-03-18 11:59:36,866 - Status of 605340a8b4f841b95295725c is RUNNING. Checking again in a minute...
2021-03-18 12:00:37,140 - Status of 605340a8b4f841b95295725c is RUNNING. Checking again in a minute...
2021-03-18 12:01:37,370 - Status of 605340a8b4f841b95295725c is RUNNING. Checking again in a minute...
2021-03-18 12:02:37,642 - Risk recalculation successful!
2021-03-18 12:02:37,642 - Closing connection to ProTego Database

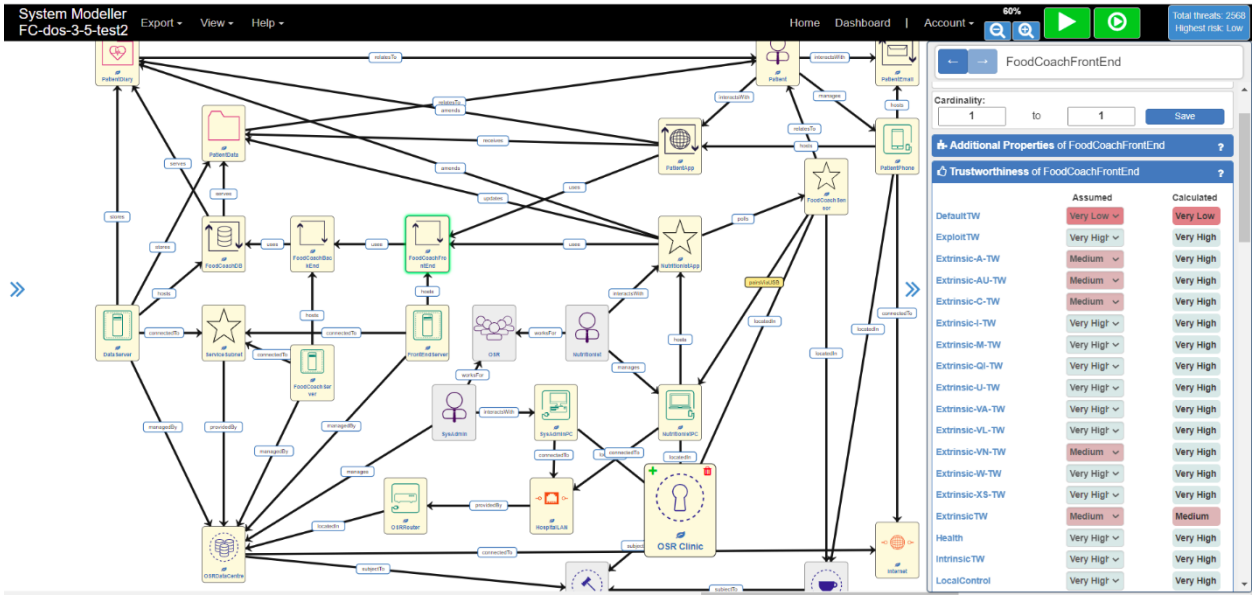
System Modeller
FC-dos-3-5-test2
Export
View
Help
Home
Dashboard
Account
60%
Total threats: 2568
Highest risk: Low
FoodCoachFrontEnd
Cardinality:
1
to
1
Save
Additional Properties of FoodCoachFrontEnd
Trustworthiness of FoodCoachFrontEnd
Assumed
Calculated
DefaultTW Very Low Very Low
ExploitTW Very High Very High
Extrinsic-A-TW Medium Very High
Extrinsic-AU-TW Medium Very High
Extrinsic-C-TW Medium Very High
Extrinsic-I-TW Very High Very High
Extrinsic-M-TW Very High Very High
Extrinsic-QI-TW Very High Very High
Extrinsic-U-TW Very High Very High
Extrinsic-VA-TW Very High Very High
Extrinsic-VL-TW Very High Very High
Extrinsic-VN-TW Medium Very High
Extrinsic-W-TW Very High Very High
Extrinsic-XS-TW Very High Very High
ExtrinsicTW Medium Medium
Health Very High Very High
IntrinsicTW Very High Very High
LocalControl Very High Very High
OSR Clinic

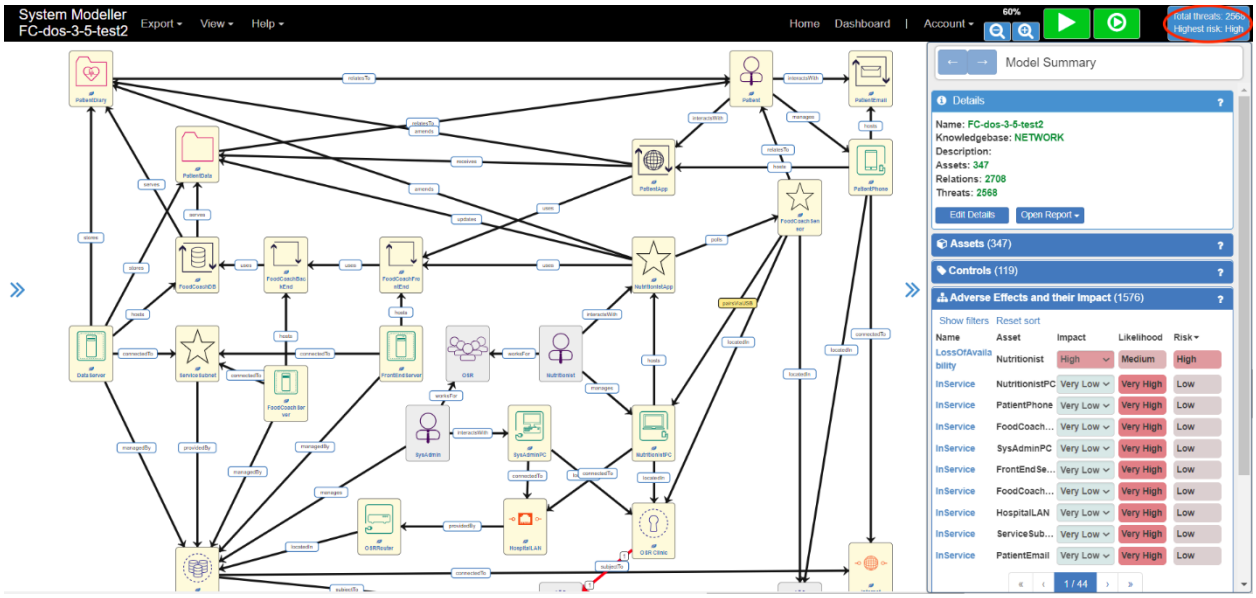
System Modeller
FC-dos-3-5-test2
Export
View
Help
Home
Dashboard
Account
60%
Total threats: 2568
Highest risk: High
Model Summary
Details
Name: FC-dos-3-5-test2
Knowledgebase: NETWORK
Description:
Assets: 347
Relations: 2708
Threats: 2568
Edit Details
Open Report
Assets (347)
Controls (119)
Adverse Effects and their Impact (1576)
Show filters
Reset sort
Name Asset Impact Likelihood Risk
LossOfAvailability Nutritionist High Medium High
InService NutritionistPC Very Low Very High Low
InService PatientPhone Very Low Very High Low
InService FoodCoach... Very Low Very High Low
InService SysAdminPC Very Low Very High Low
InService FrontEndSe... Very Low Very High Low
InService FoodCoach... Very Low Very High Low
InService HospitalLAN Very Low Very High Low
InService ServiceSub... Very Low Very High Low
InService PatientEmail Very Low Very High Low
1 / 44

ID: S25.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The tester tried to log in to the SSM with credentials that were not associated with a Security operator, and the operation failed.

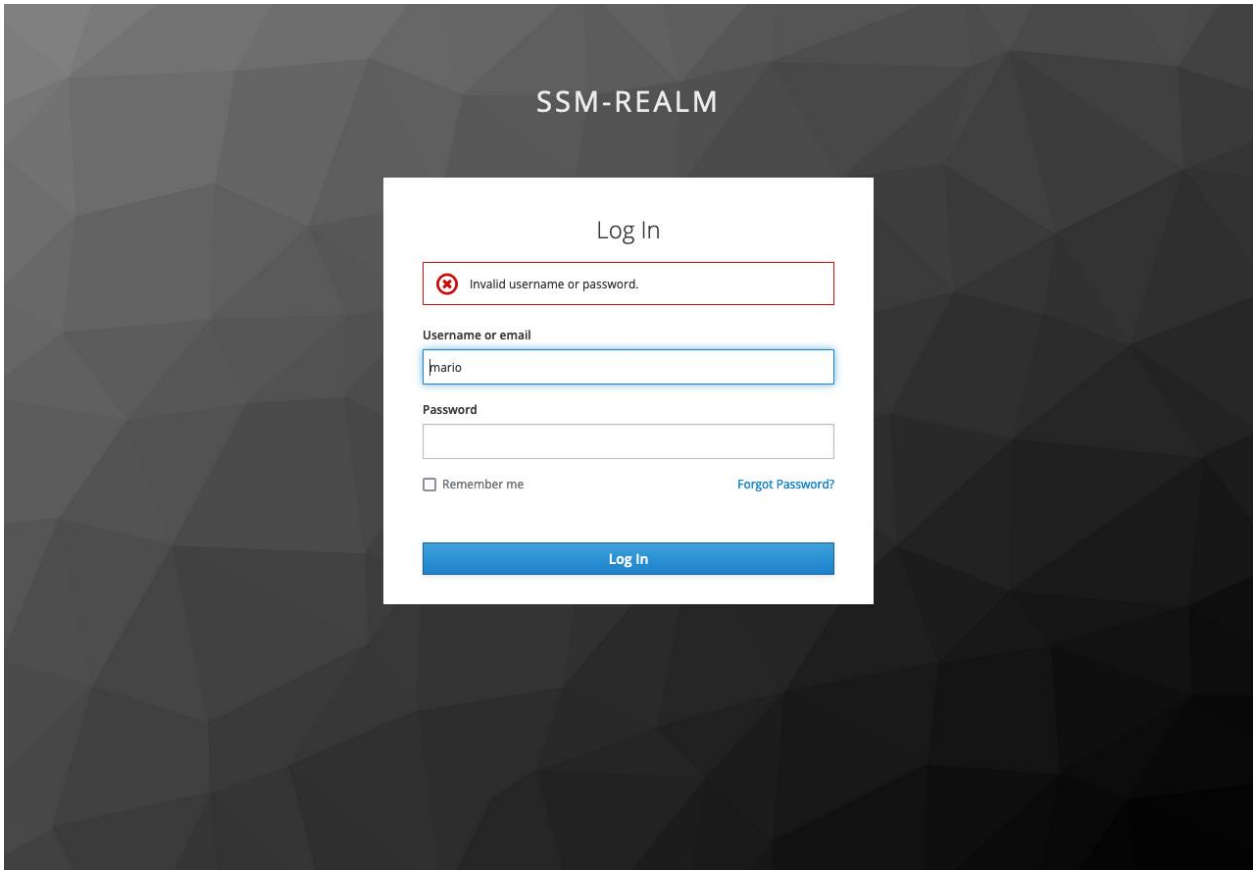

ID: S25.3

Name: First-time risk assessment is still to be conducted

Priority: Medium

Evaluation: ✓ Passed

Description: the testers are made aware that the risk assessment was not yet performed, and no models are available.



AT26: Reflect infrastructure changes

ID: S26.1
Name: Security operator reflects recent changes
Priority: High
Evaluation: ✓ Passed

Description: The tester logged into the SSM and updated the existing model (screenshot 1), thus producing a new model (screenshot 2). The updates to the model are shown in orange in the second screenshot. The risk was re-calculated by clicking the second “play” button, which resulted in: a new overall risk level (top right), and new misbehaviors (right side).

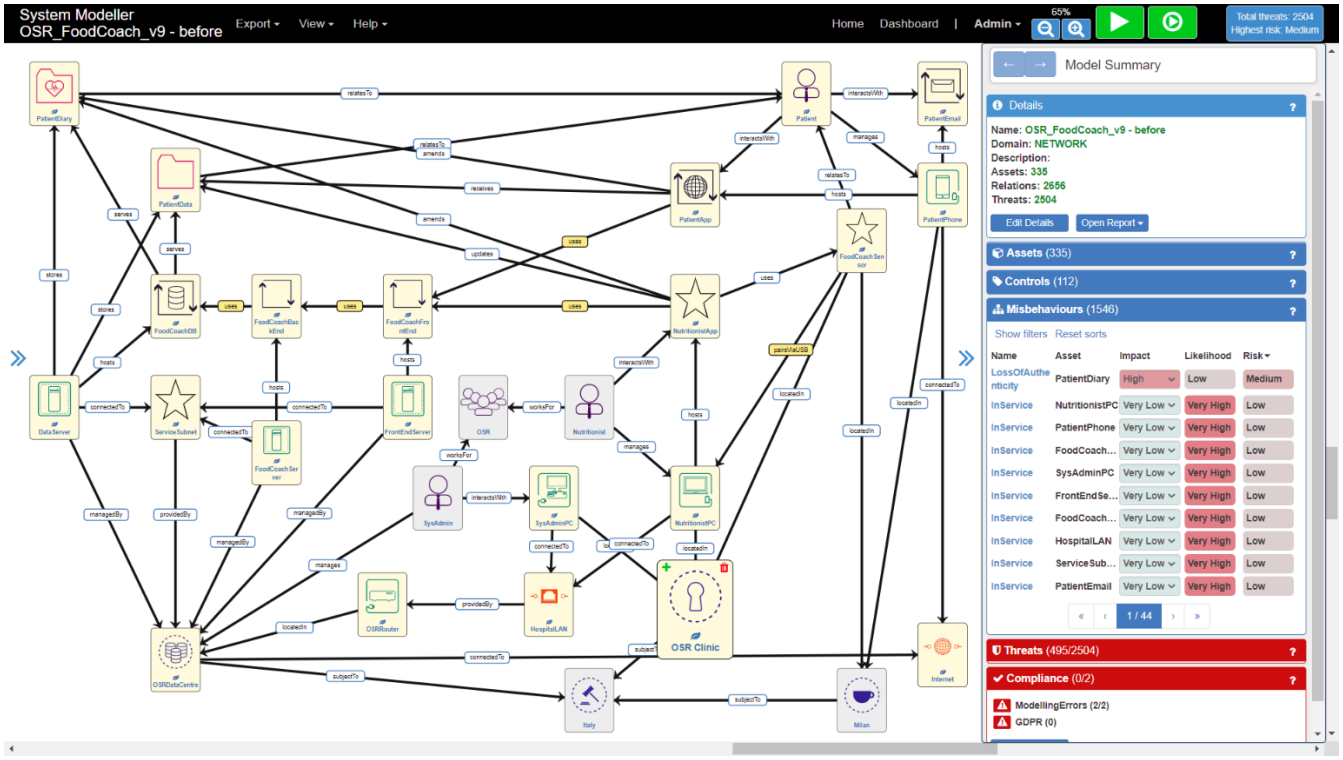

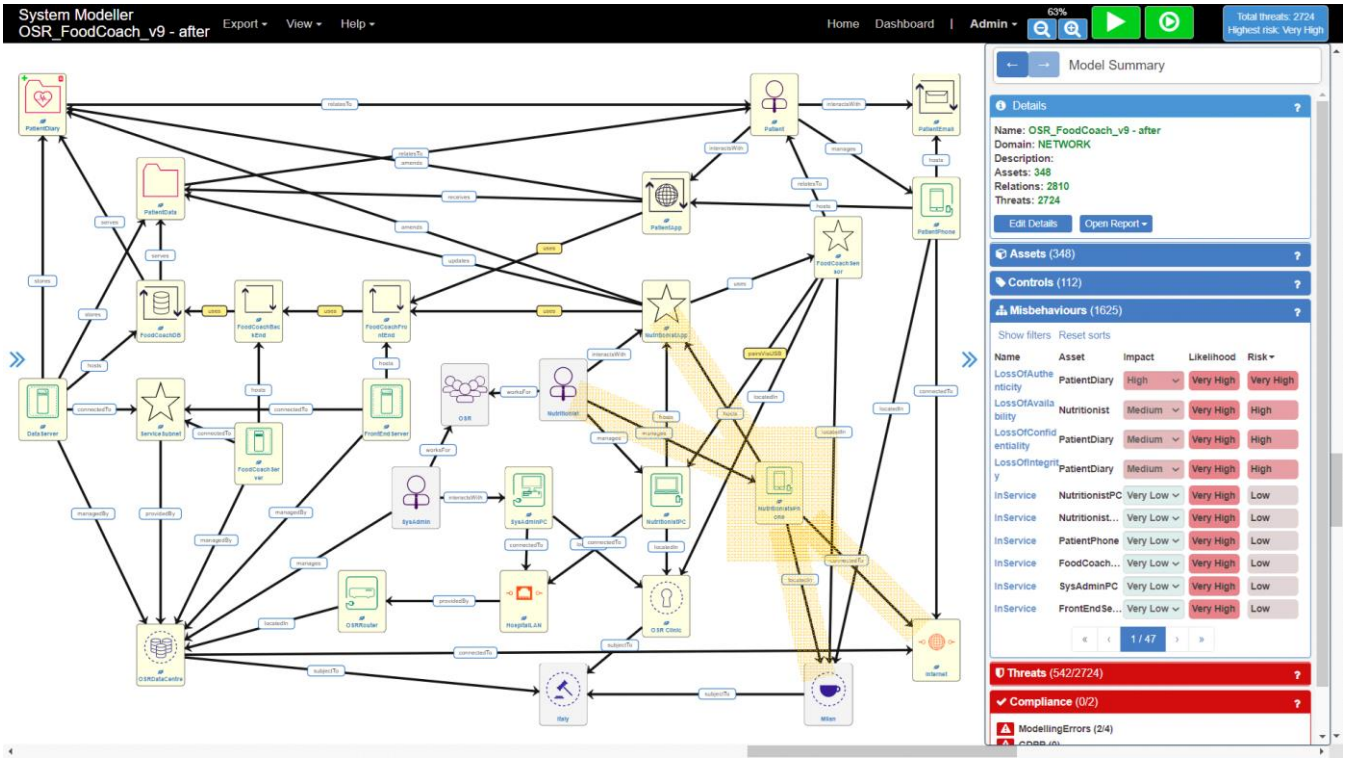
System Modeller
OSR_FoodCoach_v9 - after
Export
View
Help
Home
Dashboard
Admin
63%
Total threats: 2724
Highest risk: Very High
Model Summary
Details
Name: OSR_FoodCoach_v9 - after
Domain: NETWORK
Description:
Assets: 348
Relations: 2810
Threats: 2724
Edit Details
Open Report
Assets (348)
Controls (112)
Misbehaviours (1625)
Show filters
Reset sorts
Name
Asset
Impact
Likelihood
Risk
LossOfAuthenticity
PatientDiary
High
Very High
Very High
LossOfAvailability
Nutritionist
Medium
Very High
High
LossOfConfidentiality
PatientDiary
Medium
Very High
High
LossOfIntegrity
PatientDiary
Medium
Very High
High
InService
NutritionistPC
Very Low
Very High
Low
InService
Nutritionist...
Very Low
Very High
Low
InService
PatientPhone
Very Low
Very High
Low
InService
FoodCoach...
Very Low
Very High
Low
InService
SysAdminPC
Very Low
Very High
Low
InService
FrontEndSe...
Very Low
Very High
Low
1 / 47
Threats (542/2724)
Compliance (0/2)
ModellingErrors (2/4)

ID: S26.2

Name: Insufficient authorization

Priority: Medium

Evaluation: ✓ Passed

Description: The tester tried to log in to the SSM with credentials that were not associated with a Security operator, and the operation failed.

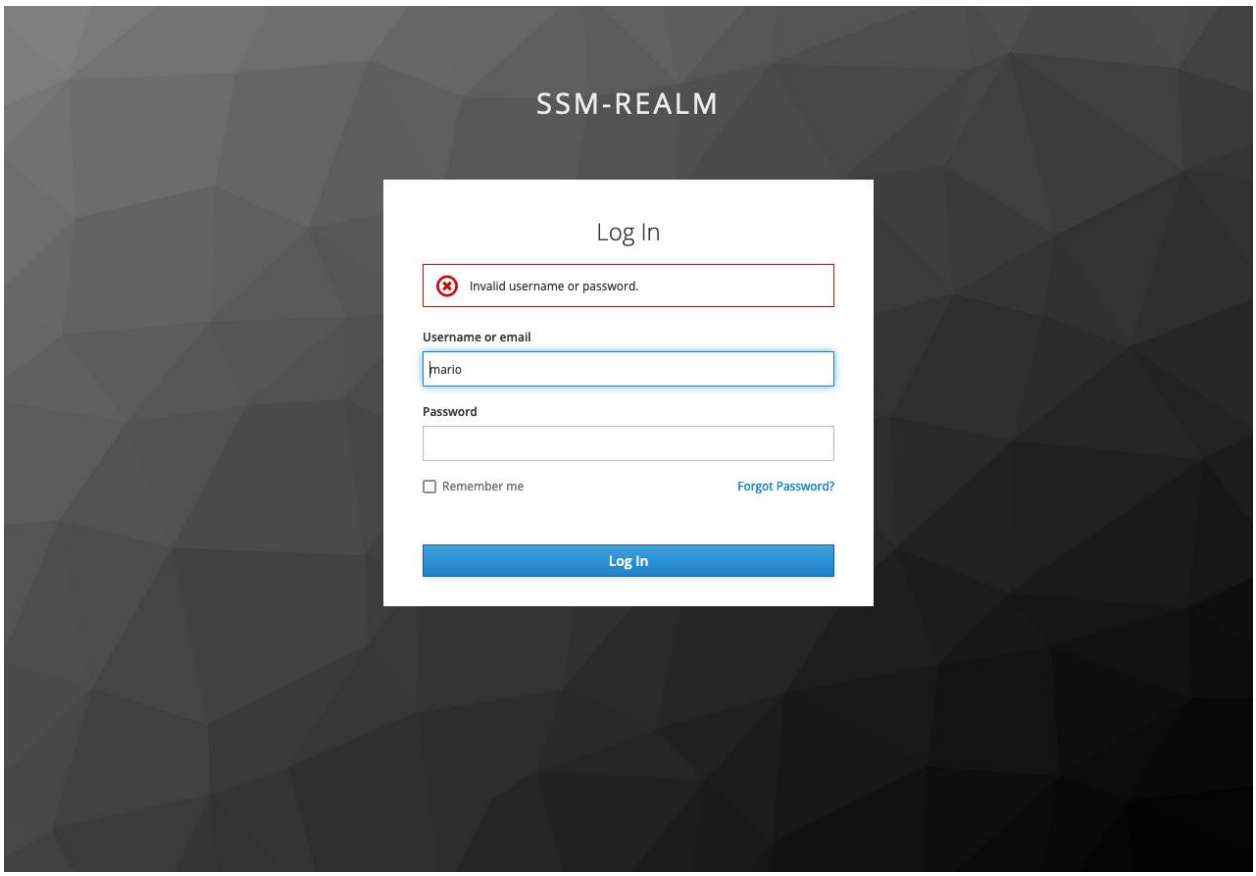

ID: S26.3

Name: First-time risk assessment is still to be conducted

Priority: Medium

Evaluation: ✓ Passed

Description: the testers are made aware that the risk assessment was not yet performed, and no models are available.



VI.2. Non-functional metrics Evaluation

Q1: How effectively is situational awareness improved?

Component: SIEM

Metric: Access Control alerts

Acceptance criterion: Alerts raised

Evaluation: ✓ Acceptable

Description: The SIEM detects alerts raised for the different security events generated by Access Control.

Component: SIEM

Metric: Continuous Authentication alerts

Acceptance criterion: Alerts raised

Evaluation: ✓ Acceptable

Description: The SIEM detects alerts raised for the different security events generated by Continuous Authentication.

Component: SIEM

Metric: Data Gateway alerts

Acceptance criterion: Alerts raised

Evaluation: ✓ Acceptable

Description: The SIEM detects alerts raised for the different security events generated by the Data Gateway.

Component: SIEM

Metric: Network Slicing alerts

Acceptance criterion: Alerts raised

Evaluation: ✓ Acceptable

Description: The SIEM detects alerts raised for the different security events generated by Network Slicing.

These four metrics show that the SIEM is able to detect alerts raised by the different components of the toolkit. Those alerts could not be detected before this functionality was implemented. Hence situational awareness is highly improved.

Q2: How efficiently does the system improve situational awareness during an attack?

Component: SIEM

Metric: Event collection rate

Acceptance criterion: 500 Events/Second per Agent

Evaluation: ✓ Acceptable

Description: The SIEM is able to collect 500 events per second per agent. Events are queued for processing as soon as the agents generate them. This number is limited by default, although it can be increased if needed.

Component: SIEM

Metric: Event processing rate

Acceptance criterion: This value depends a lot on the resources allocated (CPU and RAM), number of agents, and events per agent.

Evaluation: ✓ Acceptable

Description: Although no specific acceptance criterion was originally given, since results may vary considerably depending on the resources allocated to the testbeds, the original impression was that the processing of the events was a bit slower than expected. Nevertheless, neither a single event was dropped, nor any message was discarded as shown by the internal statistics:

```
[root@wazuh-5cbfd658b9-tqfgg /]# cat /var/ossec/var/run/ossec-analysisd.state | grep -i dropped
# Events dropped
events_dropped='0'
[root@wazuh-5cbfd658b9-tqfgg /]# cat /var/ossec/var/run/ossec-remoted.state | grep -i discarded
# Discarded messages
discarded_count='0'
```

After analyzing the behavior of the systems where the different SIEM components were running it was discovered that the performance was impacted by the storage drives of the servers. The following is the output of the `vmstat` command for the cluster worker node that runs Elasticsearch:

```
procs -----------------------memory---------------------- ---swap-- -----io---- -system-- --------cpu--------
 r  b         swpd         free         buff        cache   si   so    bi    bo   in   cs  us  sy  id  wa  st
 2  0      2034164       227484       585896       800892    0    0     0     0 8504 16091  18  17  63   0   2
 0  0      2034164       231120       585896       800892    0    0     8     4 4110 5635   9   5  79   1   5
 1  1      2034164       231616       585904       800896    0    0    28   220 5323 9300  15  11  63   9   1
 0  1      2034164       229064       585904       800896    0    0    16    56 5411 8989  10   9  52  27   3
19  0      2034164       212880       585904       800896    0    0    28     0 7365 13036  18  13  57   8   4
 0  3      2034164       230804       585904       800916    0    0     8   132 8319 17018  32  27  11  28   2
 1  2      2034164       230644       585904       800924    0    0     8   148 6676 11367  19   8  14  57   1
 0  3      2034164       230028       585924       800960    0    0    16   412 6800 12407  14  14  13  56   2
 1  6      2034164       228312       585928       800992    0    0    24   456 5932 9881  12   5  18  62   3
 4 12      2034164       221404       585928       800996    0    0    16     0 3822 7034  10   6   2  81   1
 0 16      2034164       221504       585928       801000    0    0     4     0 4510 8371  11   7   0  82   1
 2 18      2034164       221320       585928       801000    0    0     0     0 3692 6026  15   5   0  79   1
 1  9      2034164       210884       585928       801048    0    0     0    68 4346 10238  25  21   0  53   1
 0 20      2034164       221344       585928       800996    0    0    36   184 4953 8798   9   6   0  83   2
11  3      2034164       206936       585944       800980    0    0     8   116 8498 15973  19  18  22  36   4
 5  4      2034164       228716       585952       800988    0    0    64   152 10867 23264  30  33   9  22   6
 1  3      2034164       230968       585952       800996    0    0     8   136 8484 14185  15  11  21  49   4
 1  7      2034164       231268       585956       801036    0    0    12   388 6493 12152  12  15  20  51   2
 0  5      2034164       228620       585968       801012    0    0    32   284 5732 9805  15   9  11  59   6
 2  5      2034164       225800       585976       801016    0    0    28   248 8410 15230  23  14  19  39   6
 0  2      2034164       231100       585976       801004    0    0     8   128 9619 18913  22  22  19  32   5
```

The second to last column (wa – in red), indicates a very high wait time, that is, the percentage of time spent blocked while waiting for I/O to complete. The same applies to the cluster worker node where the Wazuh Manager was running, although not as high as the previous case:

```
procs -----------------------memory---------------------- ---swap-- -----io---- -system-- --------cpu--------
 r  b         swpd         free         buff        cache   si   so    bi    bo   in   cs  us  sy  id  wa  st
 5  0      1627904       297876       776896      1388148    0    0    36   196 10498 21398  20  30  46   2   2
 2  0      1627904       288864       776896      1388172    0    0     4     0 9076 20850  27  43  29   0   1
 2  0      1627904       288268       776896      1388724    0    0     8     0 8291 16499  21  40  36   0   3
 9  0      1627904       283756       776896      1388792    0    0    12     0 6789 13250  13  20  66   0   1
 9  0      1627904       291708       776896      1388800    0    0    28   168 10443 21903  21  26  33  18   3
 0  0      1627904       296664       776900      1388800    0    0    28   196 8790 16515  15  25  56   2   3
 1  1      1627904       291200       776900      1388824    0    0    12  1076 9357 19319  24  37  19  20   1
 2  1      1627904       289880       776900      1388836    0    0     8   224 9210 18774  21  41  19  18   1
 6  3      1627904       282044       776900      1388852    0    0    12   196 8415 16090  14  19  33  33   1
```

```
17  1      1627904      281456      776900      1388868    0    0    28   232 10303 22118  20  25   9  43   2
 7  1      1627904      292440      776904      1388868    0    0    12   208 6697 14282  16  23  37  22   2
 5  1      1627904      281876      776904      1388872    0    0     4     0 7568 17364  20  41  25  12   2
 3  1      1627904      291244      776904      1388888    0    0     0    16 7102 14664  23  38  23  14   2
 7  2      1627904      278808      776912      1388884    0    0     8   212 6280 12307  14  16  48  21   1
 2  2      1627904      277160      776912      1388900    0    0    20    48 6602 11874   7   9  42  41   1
 0 18      1627904      276948      776912      1388900    0    0    44   332 27885 53865  15  22   5  53   4
19  1      1627904      286348      776916      1388904    0    0    40   300 9842 25312  28  39   0  32   1
 1  1      1627904      291652      776916      1388944    0    0     8  1100 10474 26161  38  52   5   4   1
 2  4      1627904      276684      776916      1388968    0    0    12   280 7654 14772  11  23  34  31   1
12  1      1627904      282896      776916      1388960    0    0    28   232 11858 24298  21  30  18  28   2
 4  4      1627904      295268      776920      1388944    0    0    20   332 9397 18833  19  28  18  31   4
```

In both cases, after checking the drives specifications, the servers where the cluster worker nodes were running, had SATA drives with a Nominal Media Rotation Rate (NMRR) of 7200rpm:

```
host01:~$ sudo hdparm -I /dev/sda

/dev/sda:

ATA device, with non-removable media
        Model Number:       HGST HUS726020ALA610
        Serial Number:      K5HSMVJG
        Firmware Revision:  A5GNT920
        Transport:           Serial, ATA8-AST, SATA 1.0a, SATA II Extensions, SATA Rev 2.5, SATA Rev 2.6, SATA
Rev 3.0; Revision: ATA8-AST T13 Project D1697 Revision 0b
Standards:
        Used: unknown (minor revision code 0x0029)
        Supported: 9 8 7 6 5
        Likely used: 9
Configuration:
        Logical         max     current
        cylinders       16383   16383
        heads           16      16
        sectors/track   63      63
        --
        CHS current addressable sectors:    16514064
        LBA    user addressable sectors:   268435455
        LBA48  user addressable sectors:  3907029168
        Logical  Sector size:                   512 bytes
        Physical Sector size:                   512 bytes
        device size with M = 1024*1024:     1907729 MBytes
        device size with M = 1000*1000:     2000398 MBytes (2000 GB)
        cache/buffer size  = unknown
        Form Factor: 3.5 inch
        Nominal Media Rotation Rate: 7200
```

This kind of drive is evidently insufficient for a production system. At the very least SCSI drives with a NMRR of 15000rpm should be used, and if at all possible, SSD drives with the latest NVMe technology. This type of drives would significantly improve performance, since they drastically reduce the wait time for I/O.

Component: SIEM

Metric: Total log storage

Acceptance criterion: For 100 agents, 80GB/month

Evaluation: ✕ Insufficient

Description: After the period of testing needed to collect this metric, the amount of storage necessary turned out to be twice the amount originally estimated, 161GB/month. Nevertheless, although in the strictest sense this metric does not meet the original acceptance criterion, it cannot be considered as a negative point, since it includes a number of events more numerous than what is to be expected in a normal production system, since the number of tests performed increases the total amount of events generated.

Component: SIEM

Metric: Log source count

Acceptance criterion: Application, Data Gateway, Access control, Network slicing, Continuous authentication, SSM, System logs, Network devices

Evaluation: ✓ Acceptable

Description: The tests performed show that the SIEM is able to collect events from Applications, ProTego Toolkit components, Operating Systems (tests performed just with Linux and Windows systems), and Network devices (emulated Cisco systems).

Component: SIEM

Metric: Alerts triggered count

Acceptance criterion: This value can only be calculated after the SIEM has been working for a time

Evaluation: ✓ Acceptable

Description: No initial forecast or acceptance criterion was initially given to for this metric. Nevertheless, to measure it intense testing was performed on 6 different systems (Linux and Windows servers) during the course of 3 days. Within that period of time a total number of 4454 security events could be detected, which reflects the tests performed.

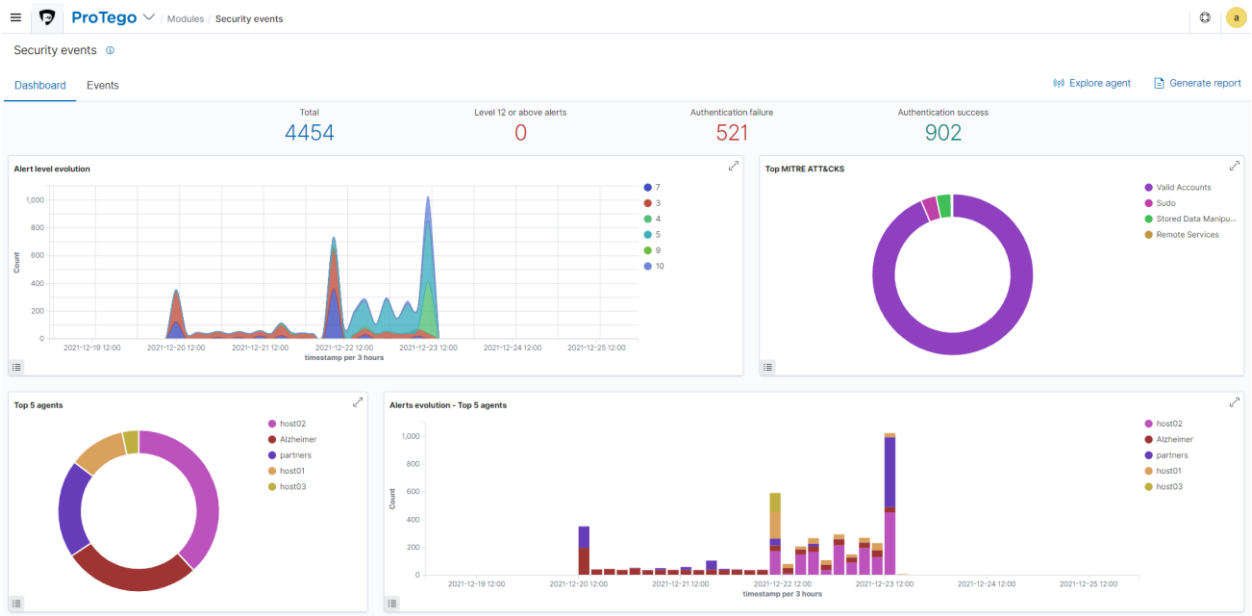

Component: SIEM

Metric: Resource usage

Acceptance criterion: CPU: 8 cores, RAM: 16 GB, Disk: 5 GB

Evaluation: ✗ Insufficient

Description: Although the goal for the SIEM was to make it work with low resources, as indicated in the acceptance criterion, to fully use all the different subsystems of the SIEM, both CPU cores and RAM had to be increased by 50% to obtain an acceptable performance. For some 100 agents the resources would have to be increased at least to double the original forecast. Besides, the storage drives, as previously indicated in the *Event processing rate* metric, require some of the commonly accepted technologies for Production systems to achieve an acceptable performance. In general, the SIEM can be considered the most demanding component of the toolkit regarding resources requirements.

Q3: Is the risk level increased as a result of new detected vulnerabilities?

Component: SIEM/SSM

Metric: Increase of risk vector level

Acceptance criterion: Overall risk level should increase as a result of the vulnerability being injected

Evaluation: ✓ Acceptable

Description: See Q4

Q4: Is the risk level reduced as a result of the implemented mitigation strategies?

Component: SSM

Metric: Reduction of risk vector level

Acceptance criterion: Overall risk level should reduce from the risk level when a vulnerability is injected as a result of the recommendations made by the SSM.

Evaluation: ✓ Acceptable

Description: An example will be used to answer both Q3 and Q4. The starting point is a risk model, and creation of such models is described in deliverables D4.2 and D4.3. The example described below uses the FoodCoach scenario from OSR as an illustration. Developments reported in D4.3 have linked the OpenVAS vulnerability scanner (https://www.openvas.org/) and the web application scanner OWASP ZAP (https://www.zaproxy.org/) with the SSM via the ProTego SIEM and Run Time Microservice. The SIEM integrates the reports from the two vulnerability scanners and the Microservice maps the vulnerabilities reported by the scanners into an internal representation used by the SSM (known as "trustworthiness attributes" – TWAs – because they represent different aspect of trustworthiness, e.g., robustness, reliability in software or motivations or knowledge in people).

The base risk model is shown in Figure VI.2-1. The service scanned by OpenVAS is the FoodCoach Front End, highlighted yellow in the figure. This is an instance of Nginx serving the FoodCoach application to the Patient via their mobile phone. This represents the normal operation state. The worst-case adverse effects (resulting from threats or vulnerabilities on assets) is at the right. The risks are ordered with the highest-level risk at the top, and it can be seen in the figure that there are 44 pages of risks in the adverse effects table, but the highest level (worst case) risks are presented at the top of the first page).

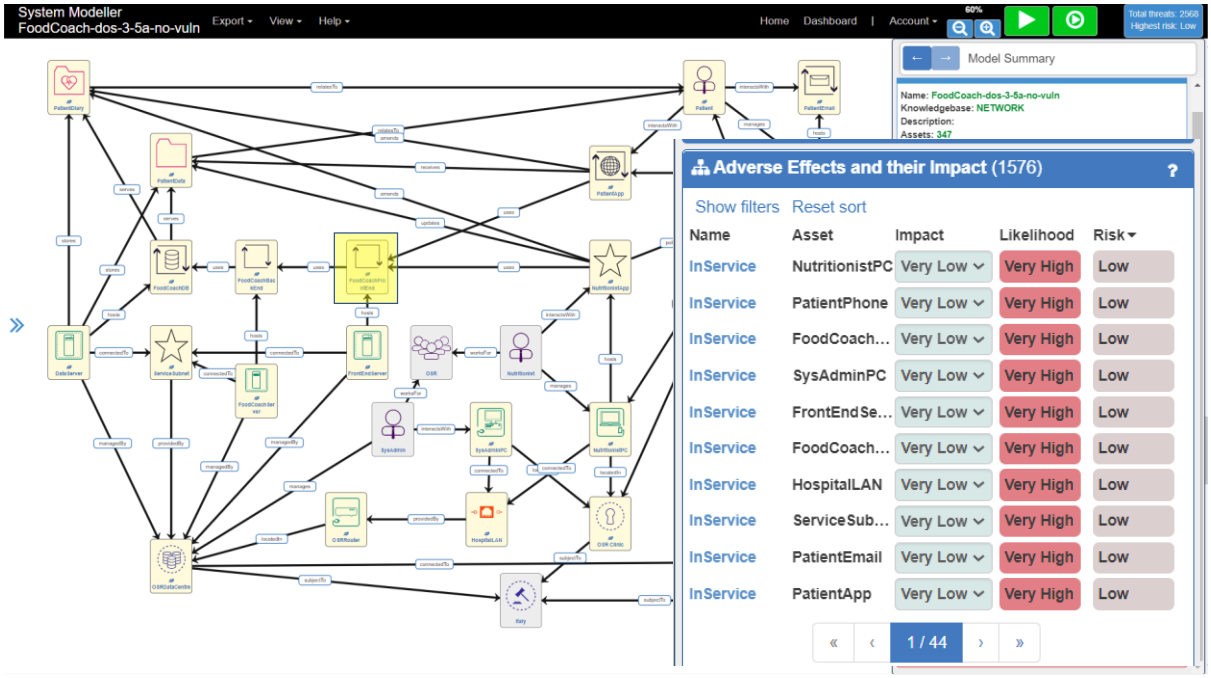

Figure VI.2-1 Normal Operation Risk Model

It can be seen that the worst-case risk is “low”, and all the risks are “InService” – these simply represent the fact that all the assets are working normally and are low level risks in this scenario.

As a result of a scan, OpenVAS reports a vulnerability, and the output is shown below.


```
{
  "cvss_version": 2,
  "identifiers": {
    "ip_address": "192.168.143.211",
    "port": "80/tcp"
  },
  "vulnerabilities": [
    {
      "name": "nginx 1.1.3 - 1.15.5 Denial of Service and Memory Disclosure
via mp4 module",
      "family": "Web Servers",
      "product": "cpe:/a:nginx:nginx:1.9.5",
      "cves": [
        {
          "CVE-2018-16845": [
            "CWE-835"
          ]
        }
      ],
      "cvss_base_vector": "AV:N/AC:M/Au:N/C:P/I:N/A:P",
      "qod": "30"
    }
  ]
}
```


This is a weakness regarding an infinite loop that could be exploited in a denial-of-service attack: *“If the loop can be influenced by an attacker, this weakness could allow attackers to consume*

excessive resources such as CPU or memory" [4]. The CVSS base vector produced ("AV:N/AC:M/Au:N/C:P/I:N/A:P") is decoded as follows[5]:

AV:N	Remote network access attack
AC:M	Access complexity Medium
Au:N	Authentication of user: No
C:P	Partial compromise of data Confidentiality
I:N	Partial compromise of data Integrity
A:P	Partial compromise of data Availability

This vulnerability is mapped into the SSM's TWAs automatically by the run time microservice, and the algorithm for this mapping is described in D4.3. A vulnerability represents a reduction in a TWA level in an asset, because the detection of the vulnerability means that the asset is less resilient or reliable and therefore more vulnerable. For this vulnerability, certain SSM TWA levels are changed as follows:

```
update OpenVAS vulnerability from dos_disc_3_5.json
list TWA changes:
TWA changes: [
  {
    "cause": "not (qi or xs) AND cvss_a: P",
    "asset_label": "FoodCoachFrontEnd",
    "twa_key": "TWAS-Extrinsic-A-TW-61d70103",
    "changed_from": "VERYHIGH",
    "changed_to": "MEDIUM"
  },
  {
    "cause": "not (qi or xs) AND cvss_c: C|P",
    "asset_label": "FoodCoachFrontEnd",
    "twa_key": "TWAS-Extrinsic-C-TW-61d70103",
    "changed_from": "VERYHIGH",
    "changed_to": "MEDIUM"
  },
  {
    "cause": "cvss_au: N",
```

[4] https://cwe.mitre.org/data/definitions/835.html

[5] https://www.first.org/cvss/v2/guide

```
    "asset_label": "FoodCoachFrontEnd",
    "twa_key": "TWAS-Extrinsic-AU-TW-61d70103",
    "changed_from": "VERYHIGH",
    "changed_to": "MEDIUM"
  },
  {
    "cause": "cvss_av: N",
    "asset_label": "FoodCoachFrontEnd",
    "twa_key": "TWAS-Extrinsic-VN-TW-61d70103",
    "changed_from": "VERYHIGH",
    "changed_to": "MEDIUM"
  }
]
```


These reductions concern the TWAs: Extrinsic-A (Availability), Extrinsic-C (Confidentiality), Extrinsic-AU (Authentication) and VN (Remote Network Access). Therefore, as a result of the vulnerability report from OpenVAS and the mapping in the run time microservice, the risk model is now more vulnerable. Running the risk calculation results in an increased level of risk, shown in Figure VI.2-2.

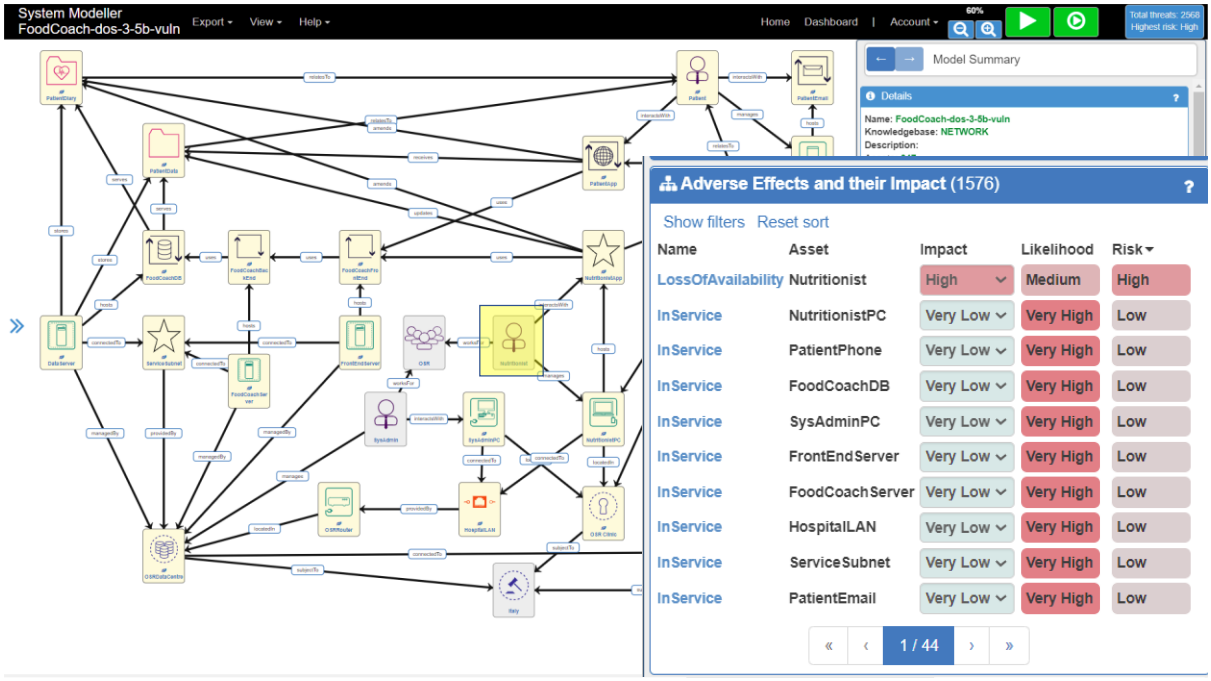

Figure VI.2-2 Increased Vulnerability Risk Model

The worst-case risk now is “high” and is a “loss of availability” at the Nutritionist (highlighted in yellow). This is because the Nutritionist’s App depends on the FoodCoach Front End being available, and if a denial-of-service attack makes the Front End unavailable, the knock-on effect is that the Nutritionist’s App is also unavailable, and the Nutritionist cannot work, rendering them unavailable. The detail of the relationships between these assets is shown in Figure VI.2-3 and Figure VI.2-4.

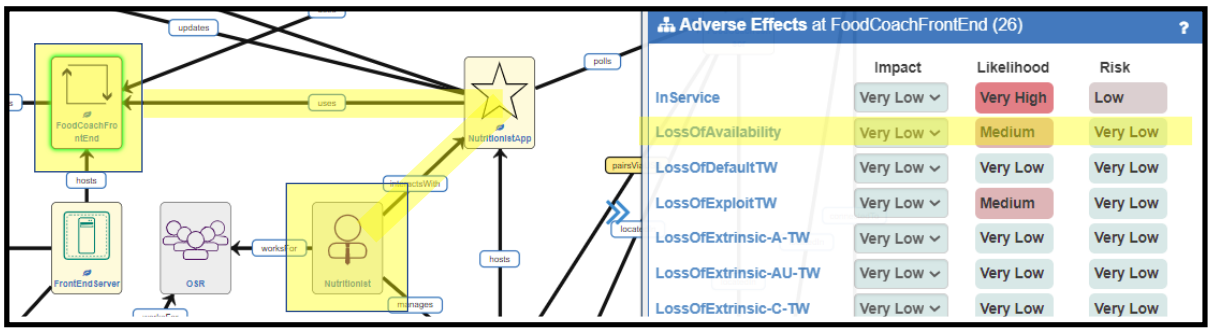

Figure VI.2-3 Dependency of Nutritionist's App on FoodCoach Front End

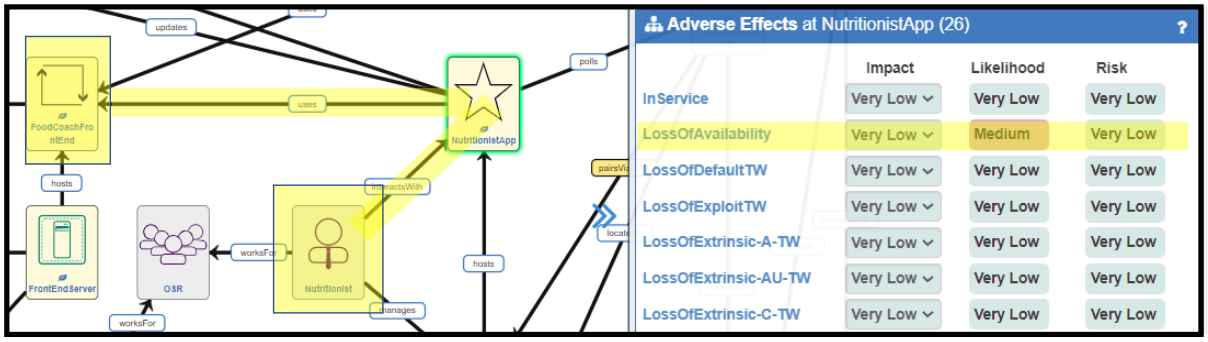

Figure VI.2-4 Dependency of Nutritionist on Nutritionist's App

Figure VI.2-3 shows the dependency of the Nutritionist's App on the FoodCoach Front End (represented by the "uses" relationship", and Figure VI.2-4 shows the dependency of the Nutritionist on their App (represented by the "interacts with" relationship). In both of these cases, if the independent entity is not available, then the dependent entity is also not available.

Looking at the "Loss of Availability" Risk at the Nutritionist, the root cause is a ***"Remote anonymous exploit on service "FoodCoachFrontEnd" from "Internet""*** – a remote attack, shown in Figure VI.2-5.

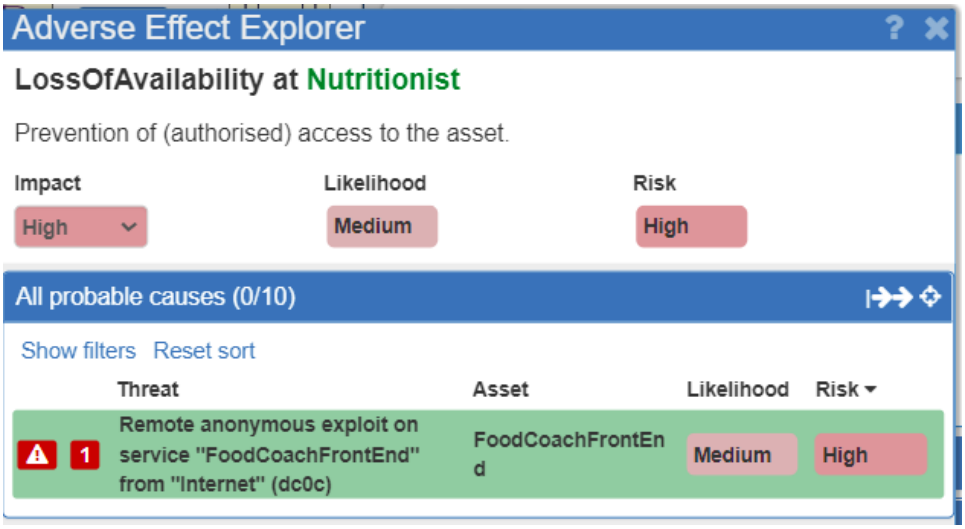

Figure VI.2-5: Root Cause Threat

The SSM recommends several approaches to addressing this threat, shown in Figure VI.2-6.



Figure VI.2-6: Control Strategy Options

The options are represented by “Control Strategies” (abbreviated to CSGs), which collect different measures applied to assets in the model. Each CSG is discussed below.

ClientAddressWhitelisting refers to a policy and practice of only accepting connections from trusted users (those on the whitelist). This is useful in some circumstances but has the disadvantages that the users must be carefully vetted and that the whitelist must be carefully managed. In the situation where there are many patients using the FoodCoach app, this approach is not practical because of the costs of vetting and management. As a result, this CSG will not be considered further.

PatchingAtProcess and ***PatchingAtProcess.Implementation*** are counterpart CSGs representing a policy to apply software patches (PatchingAtProcess) and actually applying up to date patches at runtime (PatchingAtProcess.Implementation). The policy of patching and its

implementation combined is a powerful control strategy, but its implementation relies on the availability of a patch that addresses the threat (in the case of the detected vulnerability, this refers to a patch that fixes the “infinite loop” bug in Nginx).

SuspendVulnerableServiceChannel and ***SuspendVulnerableServiceChannel.Implementation*** are counterpart CSGs representing the preparation and implementation of temporary disablement of a “service channel” – a particular route from a client to a serving process. Here the service channel in question is the path from the Patient’s phone to the FoodCoach Front End via the Internet. Suspending service channels can typically be implemented via temporary rules in firewalls, and is an effective means of blocking a remote attack threat, but at the obvious cost of the loss of the service in the channel, and can only be considered a temporary control strategy pending a more permanent solution, such as the application of a suitable software patch.

The effects of the service channel suspension and the patching will be illustrated next. Channel suspension is shown in Figure VI.2-7, and here it can be seen that the worst-case risk is “medium”, and this risk is a “Loss of Timeliness” at the Patient’s App. This means that the data in the patient’s app is out of date, and the cause of this is that the channel between the app and the FoodCoach Front End has been disabled (Figure VI.2-8 shows this in detail).

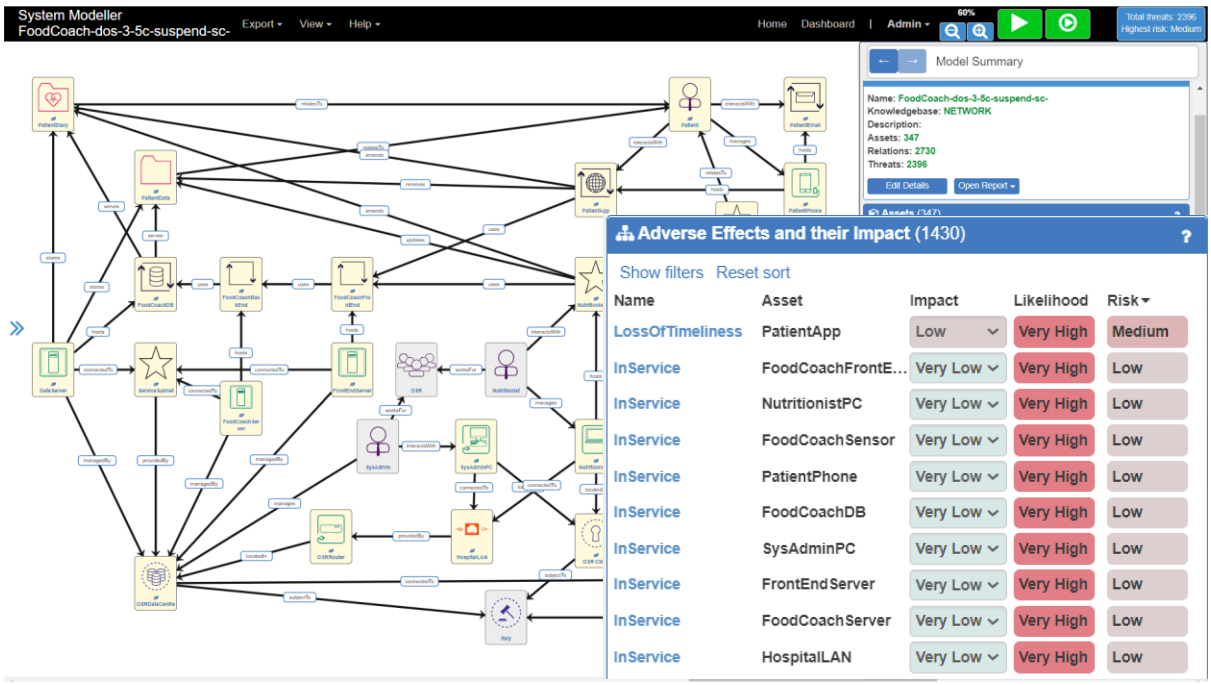

Figure VI.2-7: Channel Suspension

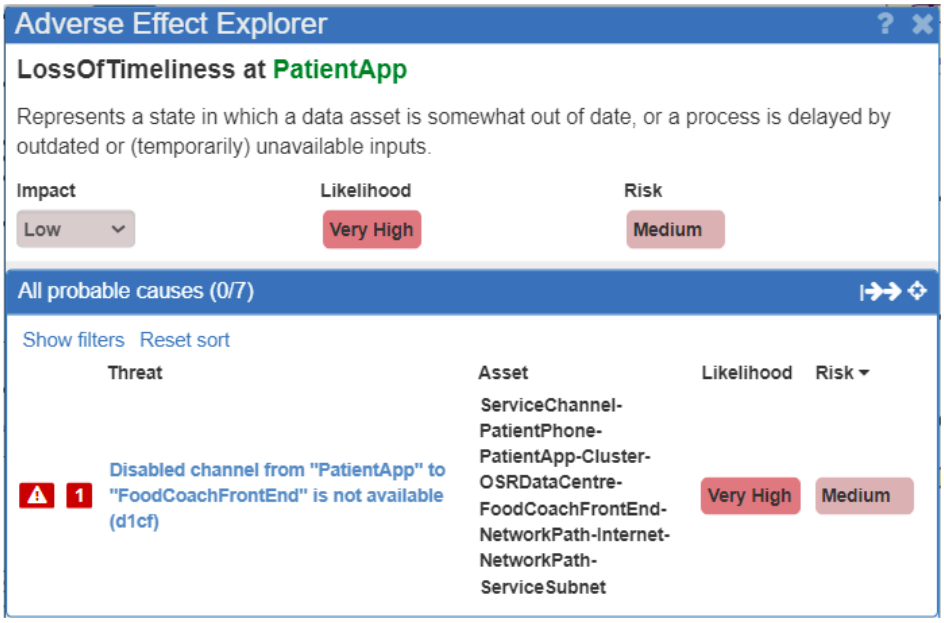

Figure VI.2-8: Detail of Loss of Timeliness at PatientApp

This example illustrates that trade-offs sometimes need to be made in controlling risk. Here, the “suspend service channel” control that addresses the vulnerability that can lead to the high risk that the Nutritionist is put completely out of action causes another lesser risk that the Patient’s app is out of date. The Loss of Timeliness risk in the Patient’s app is considered a lesser risk than the Nutritionist being unable to work because the patient can contact the Nutritionist in other ways (e.g. they may have regular appointments) and still get advice.

The effect of applying a relevant software patch is illustrated in Figure VI.2-9. Here, the worst-case risk level is “low”.

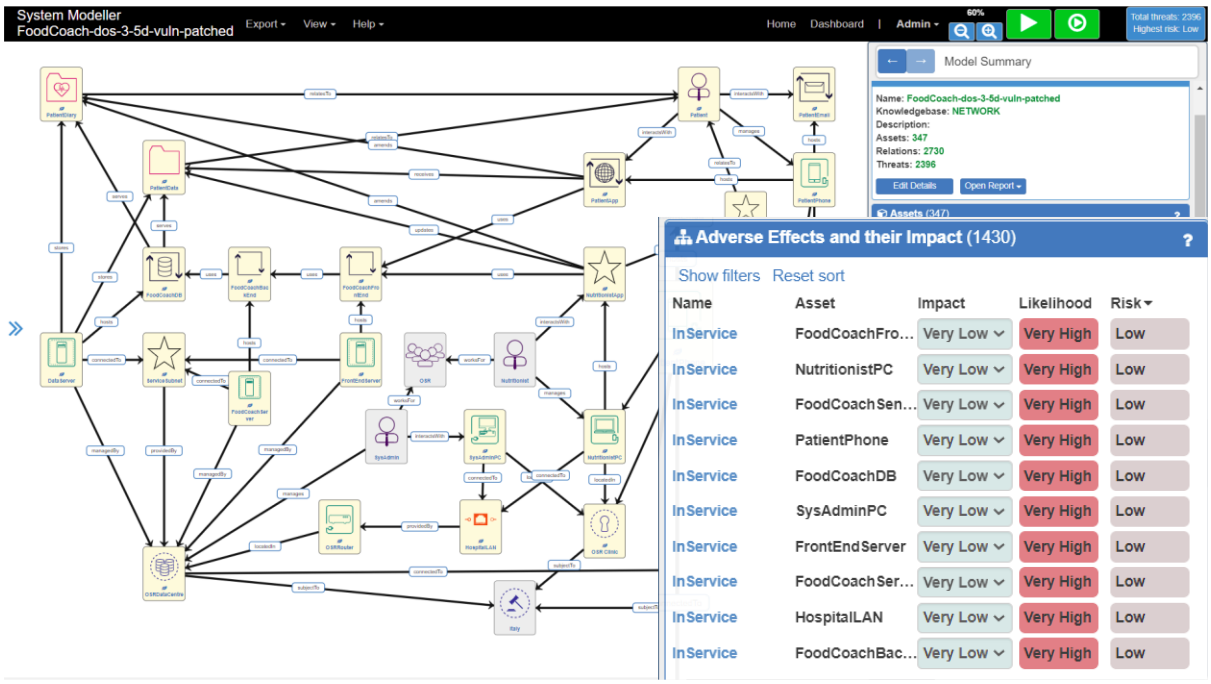

Figure VI.2-9: Software Patched Control Effect

Applying an up to date and relevant software patch is clearly more effective than disabling the service channel – here the worst-case risk is low, and the overall risk situation is the same as for normal operation (Figure VI.2-1). In Figure VI.2-9 the vulnerabilities are still present, but they are effectively controlled with no side effects. The software patching CSG is clearly the best choice, but as mentioned previously, it is dependent on a suitable patch being available.

In a real system, a suitable approach would be to combine the service channel disablement and patching CSGs, i.e., upon detection of the vulnerabilities, temporarily disable the service channel with a firewall rule, search for a suitable patch and if it available apply it, upon which the temporary disablement can be removed.

In summary, this section has illustrated how the vulnerabilities detected by vulnerability scanners (in this example OpenVAS) are reflected in the SSM risk model, causing the risk level to change (in most cases to rise to a higher level than for the status quo situation), addressing "*Q3: Does the risk level increase as a result of new detected vulnerabilities?".* This section has also illustrated the application of control strategies recommended by the SSM to lower the risk, addressing "*Q4: Does the risk level reduce as a result of the implemented mitigation strategies?*". This example has shown how this via the example of application of temporary measures (which lower the risk but at a cost) in the form of disablement, pending a permanent solution (in the form of a relevant patch being available).

Q5: How effectively are data-at-rest protected?

Component: Data Gateway

Metric: AES key size

Acceptance criterion: At least 128-bit

Result: 128-bit

Evaluation: ✓ Acceptable

Q6: How efficiently does the system perform with the data-at-rest protected?

Component: Data Gateway

Metric: Writing overhead

Acceptance criterion: max. +11% time

Result: +10.9% time

Evaluation: ✓ Acceptable

Component: Data Gateway

Metric: Reading overhead

Acceptance criterion: max. +65% time

Result: +65.8% time

Evaluation: ✓ Acceptable

Description: A careful analysis [7] was conducted to determine the reading overhead and writing overhead for using Parquet Modular Encryption. Several benchmark tests were run with different configurations, and their results are summarized in Table III.1-1. The version of the Data Gateway that is currently deployed in the ProTego toolkit is based on Java 8, uses a GCM cipher, and uses a 128-bit AES Key. In this configuration, the writing overhead and the reading overhead are +10.9% and +65.8%, respectively. It should be noted that this configuration represents the worst-case scenario. Indeed, as show in Table VI.2-1, both overheads can be greatly reduced either by using a GCM-CTR cipher, or by using a more recent version of Java (i.e., Java 11), or by using both. For example, by using both, the overheads can be reduced to just +0.9% and +2.4%, respectively.

Table VI.2-1 Parquet Modular Encryption benchmark tests

	Java8 GCM	Java8 GCM-CTR	Java11 GCM	Java11 GCM-CTR
Writing overhead	+10.9 %	+1.1 %	+2.7 %	+0.9 %
Reading overhead	+65.8 %	+14.9 %	+5 %	+2.4 %

Q7: How effectively are data-in-use protected?

Component: Access Control

Metric: False Acceptance Rate (FAR)

Acceptance criterion: FAR < 10^(-6)

Result: 1.9*10^-34

Evaluation: ✓ Acceptable

Description: The False Acceptance Rate (FAR) denotes the probability that an unauthorized user gets accepted incorrectly when requesting a specific resource. In other words, it denotes the probability that an attacker can break the authentication scheme and impersonate a legitimate user when requesting access to a resource. There are obviously multiple ways to perform such an attack, but the focus here is on attacks targeted directly at the access control framework of ProTego. Therefore, the attacker needs to forge a JWT token to perform a successful attack. As can be seen in the figure below, the algorithm to generate JWT tokens is RS256. This means that a cryptographic hash, using SHA256, is computed on the content of the token, and the result is signed using RSA with a key of size 2048 bits. According to NIST [8], this provides a security equivalent of 112-bit in a symmetric-key setting. As a result, the FAR of the system is equal to 2^-112 (= 1.9*10^-34), which is significantly below the target threshold.

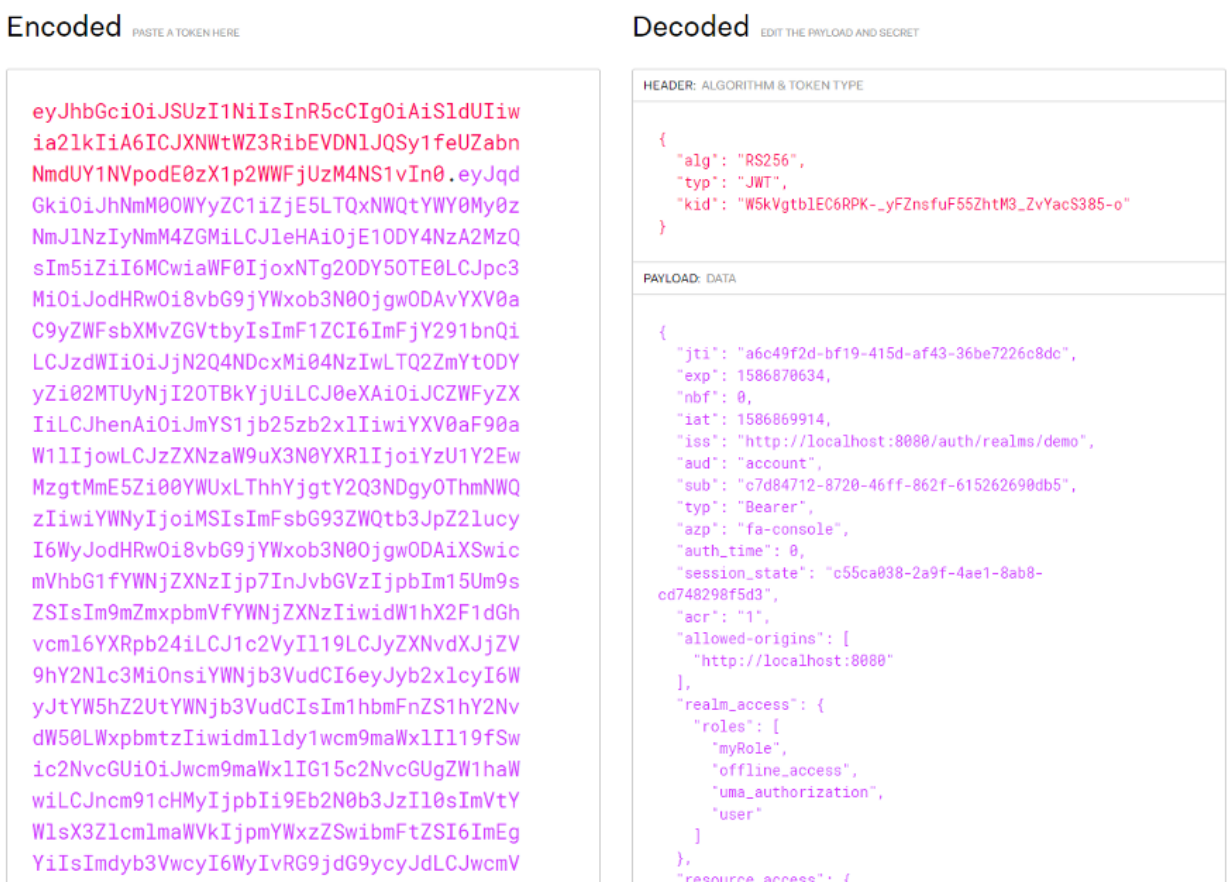

Component: Access Control

Metric: False Rejection Rate (FRR)

Acceptance criterion: FAR < 10^(-3)

Result: 0

Evaluation: ✓ Acceptable

Description: The False Rejection Rate (FRR) denotes the probability that an authorized person is incorrectly rejected when requesting a specific resource. In the system, the FRR is equal to zero. The reason for this is that by design the signature verification algorithm (i.e., the algorithm to check the correctness of the JWT token) will always succeed if it is a valid token.

Component: Access Control

Metric: Access Control type

Acceptance criterion: At least RBAC (Role Based Access Control)

Result: D-ABAC

Evaluation: ✓ Acceptable

Description: The Access Control component works on the user's attributes. These are the user ID and the role of the user (for example, patient/doctor). The component has a module called the policy decision point that contains the access control policies. This module returns the access control decision based on the above attributes and policies. Since one of the attributes is the role of the user, the access control model is at least based on RBAC.
In version 6.1 of the component, the policy decision point can provide D-ABAC, as it works on other attributes it receives from the data producer as an input file, and on attributes from the host machine environment, such as a Rego file.

Component: Access Control

Metric: Number of user roles supported

Acceptance criterion: At least two roles

Result: Doctor and Patient

Evaluation: ✓ Acceptable

Description: There are two types of users in ProTego (patient/doctor). Therefore, in the Access Control component, the access control model was defined by considering these two roles. Thus, the access control solution implemented in ProTego supports at least these two roles.

Component: Access Control

Metric: Security strength of key storage

Acceptance criterion: At least 128-bit security

Result: 256 bits

Evaluation: ✓ Acceptable

Description: The key management system of ProTego is based on Vault. Within Vault, one can specify the type of encryption key and encryption algorithm. The current specified type in the framework is as follows:

- **AES256-gcm96**: This is an authenticated encryption with associated data (AEAD) algorithm (i.e., AES-GCM) using a 256-bit key. The 96 in the name denotes the size of the nonce (not relevant for this metric).

As the name already denotes, this algorithm provides 256-bit level security. To show this metric, one can execute the cURL command:

“curl --header "X-Vault-Token: <Token>" http://127.0.0.1:8200/v1/transit/keys/<master key id>”

As a result of this command, Vault returns the response shown below. This shows again that the AES256-gcm96 is used.

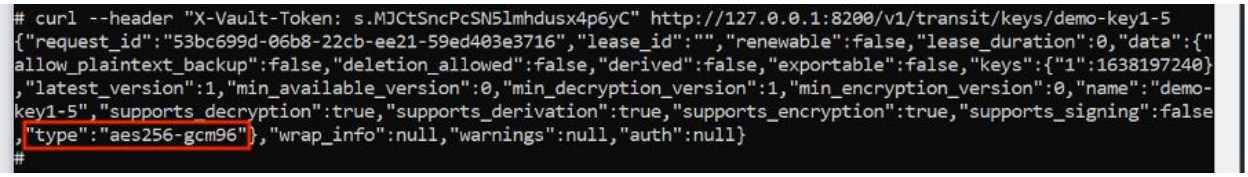

Component: Continuous Authentication

Metric: False Acceptance Rate (FAR)

Acceptance criterion: FAR < 0.5%

Result: 0.36%

Evaluation: ✓ Acceptable

Metric: False Rejection Rate (FRR)

Acceptance criterion: FRR < 2%

Result: 1.63%

Evaluation: ✓ Acceptable

Metric: Accuracy

Acceptance criterion: Accuracy > 98%

Result: 99.06%

Evaluation: ✓ Acceptable

Metric: Precision

Acceptance criterion: Precision > 95%

Result: 98.43%

Evaluation: ✓ Acceptable

Description: Participants took part in a three-session experiment. The first session consisted of typing a text of 3000 characters on their mobile phones. The second session consisted of having a casual conversation using a mobile messaging app. For the third session, a background process collected general data on app browsing and app usage during 24h. This session did not interfere with user activity on their devices. The final dataset used for the ProTego CA module included 853k events from 120 users (ranging from a minimum of 3322 events up to a maximum of 11917 events, per user). Data collected was used to feed a machine learning GBC classifier for each user that can predict identity for each single event. Results of three (FRR & Precision) or four (FAR & Accuracy) consecutive events were used to compute each metric. Values reported are the average for all 120 users.

The mathematical formula of each metric is reported below:

- FAR = FP/(FP+TN)
- FRR = FN/(FN+TP)
- Accuracy = (TP+TN)/(TP+FP+TN+FN)
- Precision = TP/(TP+FP)

where FP: False Positives, FN: False Negatives, TP: True Positives, and TN: True Negatives.

Q8: How efficiently does the system perform with data-in-use protected?

Component: Access Control

Metric: Communication overhead per access request

Acceptance criterion: Max. 1 KB

Result: 817 B

Evaluation: ✓ Acceptable

Description: This metric was obtained by using cURL, which has a built-in progress meter. The fields “Received” and “Xferd” mentioned in the output have the following meaning:

- Received: Currently downloaded number of bytes.
- Xferd: Currently uploaded number of bytes.

After executing a cURL command for a sample POST request to the access control component, the result is as follows:

```
% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 817 100 512 100 305 73142 43571 --:--:-- --:--:-- --:--:-- 132k
sending data back to the client
172.17.0.1 - - [29/Nov/2021 15:44:46] "POST /KUL/KMS/encrypt HTTP/1.1" 200 -
```

As shown in the red boxes in the figure, the values “Received” and “Xferd” are 512 bytes and 305 bytes, respectively. So, in total, we will have:

Communication overhead per access request = Received + Xferd < 1KB

Component: Access Control

Metric: Number of IAMs supported

Acceptance criterion: At least one IAM supported per hospital within ProTego

Result: Two IAMs supported

Evaluation: ✓ Acceptable

Description: In ProTego, users send their requests to the Access Control component along with the JWT token. The Access Control component only accepts JWT tokens signed by an external IAM. The Access Control implementation accepts two types of external IAMs: Keycloak and Amazon Cognito. Therefore, two IAMs are supported, one per hospital.

Q9: How effectively are data-in-transit isolated?

Component: Network Slicing

Metric: Number of slices

Acceptance criterion: Min. 3 slices

Result: 3 slices

Evaluation: ✓ Acceptable

Description: The testers successfully used Network Slicing to create a Wi-Fi network and then they configured it so that it was divided into three separate slices. The picture on the left shows the used configuration, whereas the picture on the right shows the obtained result. It's actually possible to divide the network into even more slices, if necessary.

```
root@privacygateway:/app# cat network-slicing/config.json
{
    "peer": "192.168.249.235",
    "virtual_ip": "192.0.0.1",
    "slices": [
        {
            "secure": false,
            "throughput": "40000",
            "rules": [],
            "port": 5000,
            "key": 20
        },
        {
            "secure": false,
            "throughput": "40000",
            "rules": [],
            "port": 5001,
            "key": 21
        },
        {
            "secure": true,
            "throughput": "40000",
            "psk": "swordfish",
            "rules": [{
                "protocol": "tcp",
                "address": "185.97.12.157",
                "port": "80"
            },
            {
                "protocol": "tcp",
                "address": "192.0.0.1",
                "port": "5201"
            }]
        }
    ]
}
```

```
privacy_gateway@privacygateway:~/imec-tutorial-2$ sudo ovs-vsctl show
[sudo] password for privacy_gateway:
6c009b5f-d109-46fc-84e0-3319c23377e0
    Bridge protego-slices
        Port protego-slices
            Interface protego-slices
                type: internal
        Port slice1
            Interface slice1
                type: vxlan
                options: {dst_port="5001", key="21", remote_ip="192.168.249.235"}
        Port slice2
            Interface slice2
                type: vxlan
                options: {psk=swordfish, remote_ip="192.168.249.235"}
        Port peer2
            Interface peer2
                type: patch
                options: {peer=peer1}
    Bridge protego-default
        Port slice0
            Interface slice0
                type: vxlan
                options: {dst_port="5000", key="20", remote_ip="192.168.249.235"}
        Port peer1
            Interface peer1
                type: patch
                options: {peer=peer2}
        Port protego-default
            Interface protego-default
                type: internal
    ovs_version: "2.13.3"
```

Q10: How efficiently does the system perform with the data-in-transit isolated?

Component: Network Slicing

Metric: Throughput per VPN tunnel

Acceptance criterion: Min. 100 Mbps

Result: ~300 Mbps

Evaluation: ✓ Acceptable

Description: The testers used **iperf3**, a tool that can create data streams to produce standardized performance measurements for a network. In this experiment, the testers measured the throughput from the Wi-Fi Access Point to the Privacy Gateway (Figure VI.2-10). This only includes the wired link between the Wi-Fi Access Point and the Privacy Gateway, thus excluding the link from the Wi-Fi Client to Wi-Fi Access Point, which is a wireless link with lesser throughput that would affect the measurement of the VPN tunnel itself. After running, the **iperf3** tool provides the average throughput of the measured link. As can be seen in the picture below (Figure VI.2-11), the measured throughput is over 90 Mbps, which is slightly below the required threshold of 100 Mbps, but nonetheless acceptable. The network setup at Ospedale San Raffaele includes a Switch between the Wi-Fi Access point and the Privacy Gateway, but the Switch is not strictly necessary. In fact, its presence may introduce unnecessary bottlenecks, and removing it seems to greatly improve the throughput. This, coupled with running the Privacy Gateway on bare metal (instead of Docker), results in a much higher throughput of over 300 Mbps, as demonstrated by IMEC in another experiment (Figure VI.2-12).

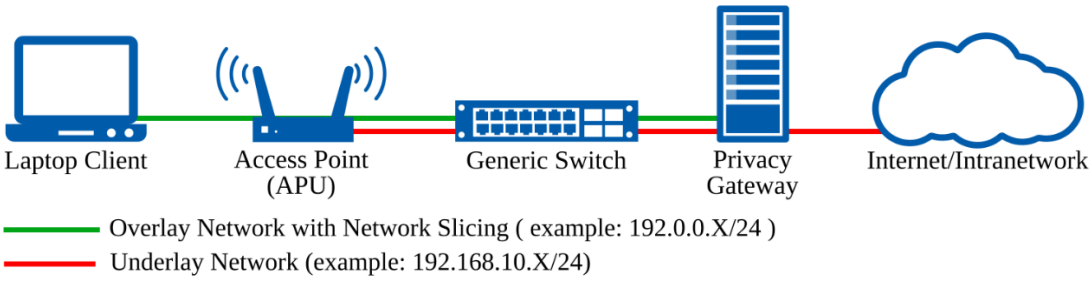

Figure VI.2-10 Network Slicing test lab

```
root@OpenWrt:~# iperf3 -c 192.0.0.1
Connecting to host 192.0.0.1, port 5201
[  5] local 192.0.0.12 port 33982 connected to 192.0.0.1 port 5201
[ ID] Interval           Transfer     Bitrate         Retr  Cwnd
[  5]   0.00-1.00   sec  11.2 MBytes  93.7 Mbits/sec    0    130 KBytes
[  5]   1.00-2.00   sec  10.7 MBytes  90.1 Mbits/sec    0    143 KBytes
[  5]   2.00-3.00   sec  10.8 MBytes  90.6 Mbits/sec    0    150 KBytes
[  5]   3.00-4.00   sec  10.8 MBytes  90.6 Mbits/sec    0    157 KBytes
[  5]   4.00-5.00   sec  10.9 MBytes  91.1 Mbits/sec    0    172 KBytes
[  5]   5.00-6.00   sec  10.9 MBytes  91.1 Mbits/sec    0    172 KBytes
[  5]   6.00-7.00   sec  10.9 MBytes  91.1 Mbits/sec    0    172 KBytes
[  5]   7.00-8.00   sec  10.9 MBytes  91.6 Mbits/sec    0    229 KBytes
```

Figure VI.2-11 Network performance test 1

```
root@OpenWrt:/# iperf3 -c 192.0.0.1
Connecting to host 192.0.0.1, port 5201
[  5] local 192.0.0.2 port 40744 connected to 192.0.0.1 port 5201
[ ID] Interval           Transfer     Bitrate         Retr  Cwnd
[  5]   0.00-1.01   sec  36.5 MBytes   301 Mbits/sec    0    203 KBytes
[  5]   1.01-2.00   sec  37.5 MBytes   318 Mbits/sec    0    203 KBytes
[  5]   2.00-3.02   sec  38.8 MBytes   319 Mbits/sec    0    203 KBytes
[  5]   3.02-4.03   sec  38.8 MBytes   322 Mbits/sec    0    203 KBytes
[  5]   4.03-5.02   sec  37.7 MBytes   320 Mbits/sec    0    255 KBytes
[  5]   5.02-6.03   sec  38.8 MBytes   323 Mbits/sec    0    255 KBytes
[  5]   6.03-7.03   sec  38.8 MBytes   323 Mbits/sec    0    255 KBytes
[  5]   7.03-8.00   sec  37.5 MBytes   323 Mbits/sec    0    255 KBytes
[  5]   8.00-9.01   sec  38.8 MBytes   323 Mbits/sec    0    255 KBytes
[  5]   9.01-10.02  sec  38.8 MBytes   323 Mbits/sec    0    255 KBytes
- - - - - - - - - - - - - - - - - - - - - - - - -
[ ID] Interval           Transfer     Bitrate         Retr
[  5]   0.00-10.02  sec   382 MBytes   320 Mbits/sec    0             sender
[  5]   0.00-10.06  sec   382 MBytes   318 Mbits/sec                  receiver

iperf Done.
```

Figure VI.2-12 Network performance test 2

Component: Network Slicing

Metric: Latency

Acceptance criterion: Max. 5 ms

Result: ~4.3 ms

Evaluation: ✓ Acceptable

Component: Network Slicing

Metric: Packet loss

Acceptance criterion: Max. 0.001%

Result: 0%

Evaluation: ✓ Acceptable

Description: The testers used the **ping** command-line utility to send echo request packets from the Wi-Fi client (i.e., the tester's laptop) to the Privacy Gateway via Internet Control Message Protocol (ICMP). The statistics of the performed experiment can be easily inferred from the program's output, which is shown in the picture below (Figure VI.2-13). Most importantly, the packet loss is 0% and the average latency is around 4.3 milliseconds, which is within the expected threshold of 5 milliseconds. The output of the program is truncated in the picture for the sake of brevity, but the results are wholly reported in the cumulative percentage graph below (Figure VI.2-14) to offer a more comprehensive view of the results. As shown in the graph, out of 150 packets, 113 packets (75%) have a latency between 2 and 4 milliseconds, while 21 packets (14%) have a latency between 4 and 6 (excluded) milliseconds. Only 16 packets (10%) have a latency that exceeds the threshold, but these packets are overall sporadic and do not push the average latency beyond the acceptable value of 5 milliseconds. These sudden spikes are rare occurrences but typical of wireless networks, and they can be attributed to wireless interference, which occurs when the Wi-Fi signal of the router is disrupted by any other device that emits an electro-magnetic signal in its proximity.

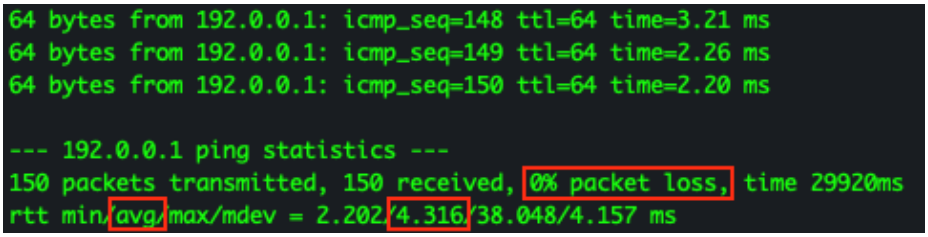

Figure VI.2-13 Ping test

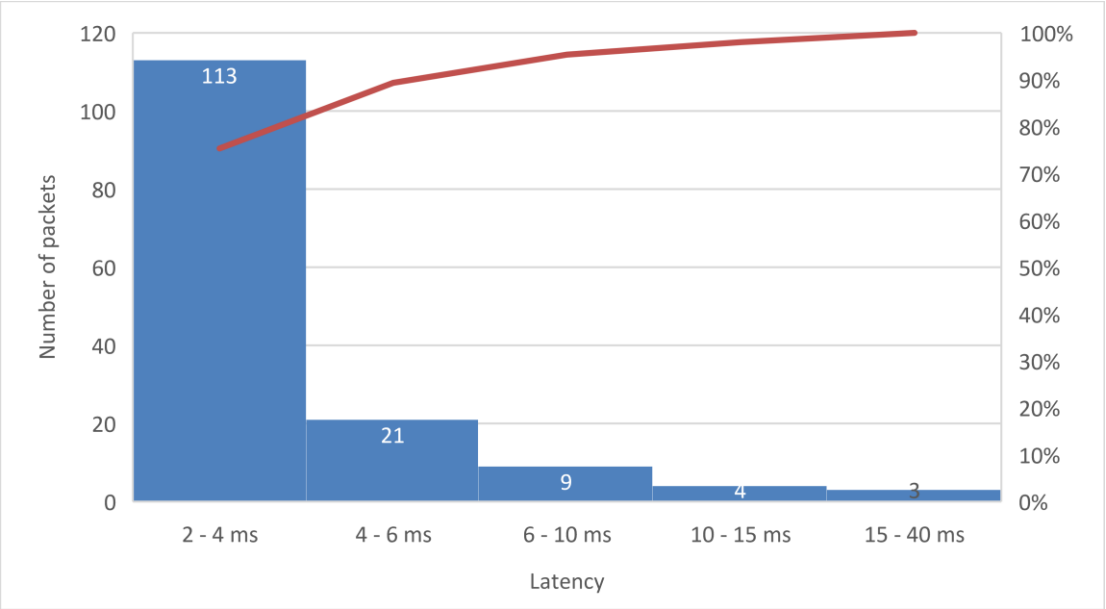

Figure VI.2-14 Ping test statistics

Component: Network Slicing
Metric: Throughput

Acceptance criterion: Min. 10 Mbps

Result: ~14 Mbps

Evaluation: ✓ Acceptable

Description: Once again, the testers used the **iperf3** tool to measure the throughput. Unlike the "Throughput per VPN" metric, in this case the experiment was run from the Wi-Fi Client to the Privacy Gateway, thus involving the wireless channel which includes the performance isolation implemented by Network Slicing. As can be seen in the picture below (Figure VI.2-15), the measured throughput is over 14 Mbps, which is well above the required threshold of 10 Mbps.

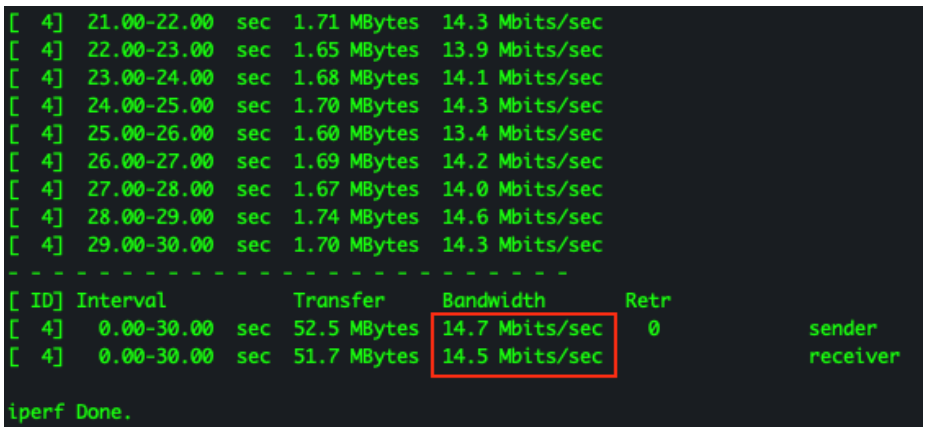

Figure VI.2-15 Network Slicing throughput test

Component: Network Slicing

Metric: VPN additional delay

Acceptance criterion: Max. -0.7 ms

Result: +0.2 ms

Evaluation: ✓ Acceptable

Description: Once again, the testers used the **ping** command-line utility to send echo request packets to the Privacy Gateway via Internet Control Message Protocol (ICMP). ICMP traffic is generally not encrypted, but in this case, encryption was purposely enabled to verify whether it resulted in increased latency. As can be seen in picture (Figure VI.2-16), the transmitted packets have an average latency of around 4.1 milliseconds, whereas the average latency for unencrypted data was around 4.3 milliseconds, as shown in the previous experiment (Figure VI.2-13). This obviously does not mean that encrypting traffic makes it any faster, but rather, that the impact of encryption on latency is practically negligible. Rather, latency appears to be more affected by other unpredictable factors, most importantly wireless interference.

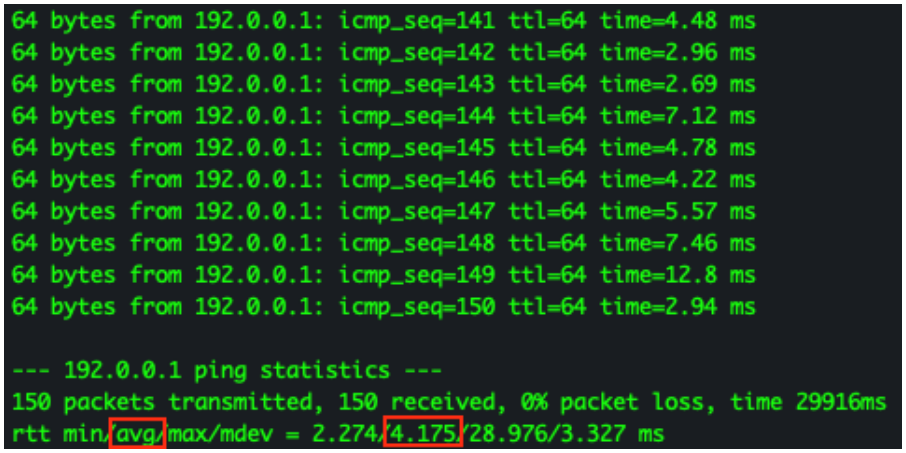

Figure VI.2-16 Ping test encrypted

VI.3. Integration Toolkit New Features

This section describes the work and final status of the Integration Toolkit and integrated components after installing the new features Istio and Storage and updating components to integrate with those new features.

VI.3.1. Istio

The first new feature for the Integration Toolkit is Istio, a service mesh technology mainly used in ProTego for traffic management and some security features, as described in D6.3. This new feature will also allow the ProTego components to be exposed via URL names not IP addresses.

After installing Istio, all components' charts from WP4 and WP5 have been updated with Istio Gateway and Istio Virtual Services resources.

This is an example of an Istio resource added for the SIEM where the Kibana port 5601 is being exposed through port 80:

```
{{- if .Values.istio.enabled }}
apiVersion: networking.istio.io/v1alpha3
kind: Gateway
metadata:
  name: kibana-gateway
spec:
  selector:
    istio: ingressgateway
  servers:
  - port:
      number: 80
      name: kibana-http
      protocol: HTTP
    hosts:
    - {{ .Values.istio.url }}
---
apiVersion: networking.istio.io/v1alpha3
kind: VirtualService
metadata:
  name: kibana-virtualservice
spec:
  hosts:
  - {{ .Values.istio.url }}
  gateways:
  - kibana-gateway
  http:
  - match:
    - uri:
```

```
            prefix: /
        route:
        - destination:
            host: kibana
            port:
              number: 5601
    {{- end }}
```

Integration with Istio has been different for some components. Some issues were found with the SIEM and SSM integration with the Istio sidecar, for several reasons depending on the components, thus it was decided not not to use it for those components, and just use the ingress for them. This might be fixed in future work. Table VI.3-1 represents the final status of the Istio integration:

Table VI.3-1 Istio integration by toolkit component

WP	Component	Istio Ingress	Istio Sidecar
WP4	SSM	✓	
WP4	SIEM	✓	
WP5	Data Gateway	✓	✓
WP5	Access Control Framework	✓	✓
WP5	Continuous Authentication	✓	✓
WP5	Network Slicing	✓	✓

After updating the charts, some additional work had to be done in order to expose and load balance the components. There are some differences in this regard between the OSR and Marina Salud deployments due to OSR being and OnPremise and Marina Salud an OnCloud deployment.

At OSR, an external nginx reverse proxy had to be deployed in order to load balance to the nodes of the Kubernetes cluster. The Istio ingress is exposed via NodePorts at the cluster, then the external nginx will redirect all traffic to the Kubernetes nodes and corresponding port. Figure VI.3-1 displays the traffic flow from nginx to the Kubernetes cluster.

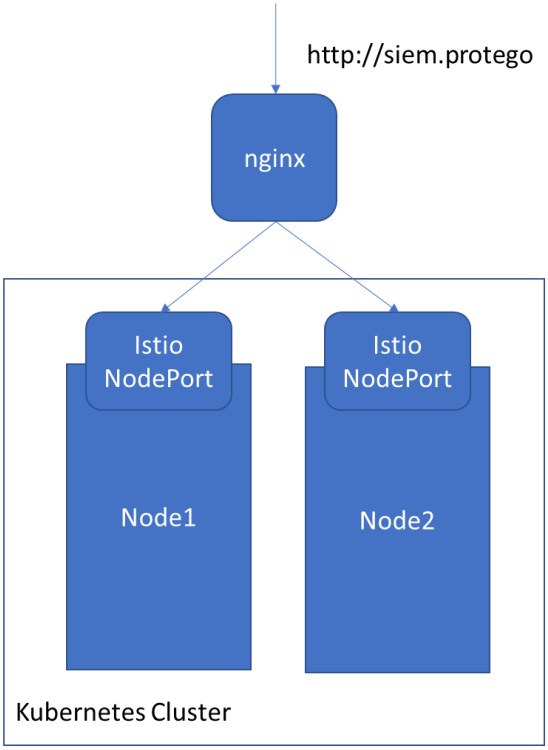

Figure VI.3-1 Integration Toolkit Traffic Flow

At Marina Salud, the nginx was not needed since Amazon ELB (LoadBalancer) was used. It was configured to perform L4 redirections for the specific ports being exposed. The Architecture is the same as for OSR just replacing the nginx by the ELB.

VI.3.2. Storage

Regarding the storage, ss described in D6.3, there are to main issues when using local storage. If a Pod, with a local persistent volume, crashes, Kubernetes may choose to deploy it to a different node. In that case, data is not moved to the new node, thus resulting in data loss from the application. Another issue happens when different Pods need to share a persistent volume. In that case, both Pods need to be in the same node to access the same data in the volume. To address those issues a different storage technology has been used in the Integration Toolkit.

Regarding the first issue, Longhorn has been deployed in the cluster. Longhorn is a highly available persistent storage for Kubernetes, which uses the local filesystem of the nodes and provides replication and backups. It is deployed on a Kubernetes cluster as a regular application using a Helm chart. It is used by the application by selecting the specific StorageClass in a Persistent Volume Claim. This allows Pods to be re-scheduled to different nodes of the cluster maintaining data availability. Charts were updated for every component using this technology, changing the persistent volume claim to use the specific Longhorn StorageClass.

This is an example of a persistent volume claim for the Continuous Authentication EDR Mongo database using Longhorn. Notice the "*storageClassName*" field set to "*longhorn*":

```
{{- if eq .Values.volumes.type "longhorn" }}
{{ if not .Values.mongo.volumes.data.existingClaim }}
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
```

```
  name: {{ template "ca-edr.fullname" .}}-mongo-data-pvc
  namespace: {{ .Release.Namespace }}
spec:
  {{- if eq .Values.volumes.type "longhorn" }}
  storageClassName: longhorn
  accessModes:
    - ReadWriteMany
  {{- end }}
  resources:
    requests:
      storage: {{ .Values.mongo.volumes.data.pvc.capacity }}
{{- end }}
{{- end }}
```

Regarding the sharing of storage, NFS is used as already described in D6.3. The corresponding charts were also updated, in this case, changing the persistent volume to use the "*nfs*" option provided by Kubernetes out of the box, which allows Kubernetes to mount NFS volumes on the Pods.

This is an example of the Data Gateway persistent volume to use NFS. Notice the "*nfs*" section where the NFS server address and volume to be mounted are set:

```
{{- if .Values.persistence.enabled -}}
apiVersion: v1
kind: PersistentVolume
metadata:
  namespace: {{ .Release.Namespace }}
  name: {{ .Values.app.name }}-pv-volume
  labels:
    type: local
spec:
  storageClassName: manual
  capacity:
    storage: {{ .Values.persistence.storage }}
  accessModes:
    - {{ .Values.persistence.accessmode }}
  {{- if eq .Values.volumes.type "local" }}
  hostPath:
    path: "{{ .Values.volumes.basepath }}/els-fhir"
    type: DirectoryOrCreate
  {{- else }}
  nfs:
    server: {{ .Values.volumes.nfs_server }}
```

```
      path: "{{ .Values.volumes.nfs_basepath }}/els-fhir"
    {{- end }}

  # hostPath:
  #   path: {{ .Values.persistence.path | quote }}
{{- end }}
```

Similar to Istio, the storage integration has been also different depending on the components, using Longhorn or NFS depending on the specific application use case. In addition, there were some issues with the SIEM Vulnerabilities components, due to some restrictions in terms of database communication that prevented from using any of those technologies, so it was configured to use local filesystem. This might also be investigated in future work. Table VI.3-2 represents the final status of the storage integration:

Table VI.3-2 Storage integration

WP	Component	NFS	Longhorn
WP4	SSM		✓
WP4	SIEM		✓
WP5	Data Gateway	✓	
WP5	Access Control Framework	✓	
WP5	Continuous Authentication		✓
WP5	Network Slicing		✓

VII. REFERENCES AND INTERNET LINKS

[1] K. Wiegers, and J. Beatty, “Software requirements,” Pearson Education, 2013.

[2] S. Covey, “The 7 Habits of Highly Effective People,” New York: Free Press, 2004.

[3] K. Nicieja, “Writing Great Specifications: Using Specification by Example and Gherkin,” Manning Publications, 2017.

[4] A. William, and T.Tullis, “Measuring the user experience: collecting, analyzing, and presenting usability metrics,” Newnes, 2013.

[5] J. R. Lewis, “Psychometric evaluation of an after-scenario questionnaire for computer usability studies: the ASQ.” SIGCHI Bull. 23, 1, 1991.

[6] https://xd.adobe.com/ideas/process/user-testing/sus-system-usability-scale-ux/

[7] https://medium.com/@tomersolomon/test-driving-parquet-encryption-3d5319f5bc22

[8] https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-57pt1r5.pdf


ProTego