Capabilities and Use of Cortex as part of the DYNAMO project

Kennedy Kalonji

2022 Laurea

Laurea University of Applied Sciences

Capabilities and Use of Cortex as part of the DYNAMO project

Kennedy Kalonji
Business Information Technology
Thesis
May 2022

Laurea University of Applied Sciences **Abstract**

Degree Programme in Business Information Technology, Cybersecurity

Bachelor of Business Administration

Kennedy Kalonji

Capabilities and Use of Cortex as part of the DYNAMO project

Year 2022 Number of pages 54

According to the 2022 Cyber Threat Report from SonicWall, the volume of ransomware attacks in 2021 grew by 105% from 2020 and was multiplied by three compared to 2019. This alarming signal demonstrates the need for more cybersecurity resilience and business continuity planning for companies worldwide. The greatest challenges are for those companies operating in critical sectors such as energy, healthcare, or maritime transportation for which cyber-attacks can have disastrous implications. The DYNAMO project was created to address this issue. A specific task of the project, with Laurea University of Applied Sciences as lead, involves the development of a module allowing the storing and sharing of threat information within the community of DYNAMO participants in an environment that is trustworthy.

My work consists in shedding light on the open-source tool Cortex chosen for developing the module by listing the capabilities of Cortex and replacing them in the context of the Computer Security Incident Response framework and by demonstrating the use of Cortex through the creation of a test environment. Meeting these objectives will allow the project to ensure that they make the best use of Cortex and that expectations related to Cortex in the development of the module are met.

Keywords: threat information sharing, cyber security, open-source tool, observable analysis

Laurea-ammattikorkeakoulu **Tiivistelmä**

Tietojenkäsittely

Tradenomi (AMK)

Kennedy Kalonji

Cortexin ominaisuudet ja käyttö osana DYNAMO-projektia

Vuosi 2022 Sivumäärä 54

SonicWallin vuoden 2022 Cyber Threat Reportin mukaan kiristysohjelmahyökkäysten määrä vuonna 2021 kasvoi 105 % vuodesta 2020 ja kolminkertaistui vuoteen 2019 verrattuna. Tämä hälyttävä signaali osoittaa, että yrityksille maailmanlaajuisesti tarvitaan lisää kyberturvallisuuden sietokykyä ja liiketoiminnan jatkuvuuden suunnittelua. Suurimmat haasteet ovat niillä yrityksillä, jotka toimivat kriittisillä aloilla, kuten energia-, terveydenhuolto tai merikuljetus, joille kyberhyökkäyksillä voi olla tuhoisia seurauksia. DYNAMO-projekti luotiin tämän ongelman ratkaisemiseksi. Hankkeen erityistehtävä Laurea-ammattikorkeakoulun johdolla on moduulin kehittäminen, joka mahdollistaa uhkatietojen tallentamisen ja jakamisen DYNAMOn osallistujien yhteisön sisällä luotettavassa ympäristössä.

Työni koostuu moduulin kehittämiseen valitun avoimen lähdekoodin työkalun Cortexin valaisemisesta listaamalla Cortexin ominaisuudet ja korvaamalla ne Computer Security Incident Response -kehyksen yhteydessä sekä esittelemällä Cortexin käyttöä testin avulla. ympäristöön. Näiden tavoitteiden saavuttaminen mahdollistaa sen, että hanke varmistaa, että Cortexia hyödynnetään parhaalla mahdollisella tavalla ja että Cortexiin liittyvät odotukset moduulin kehittämisessä täyttyvät.

Asiasanat: uhkatietojen jakaminen, kyberturvallisuus, avoimen lähdekoodin työkalu, havaittava analyysi

Contents

1 Introduction

The DYNAMO project was created to respond to the need of companies from critical sectors (Healthcare, Energy and Maritime Transportation) for the development of a methodology and a solution that would assist them in their cybersecurity resilience and business continuity planning. Laurea is part of the project and is in charge of the task 4.2. which consists in developing a module for storing and sharing threat information within the community of DYNAMO participants in an environment that is trustworthy. The open-source tools MISP, TheHive and Cortex have been identified as tools on which the project will be based. To this end, this paper will focus on the open-source tool Cortex which a powerful tool used to analyse observables and generate active responses based on defined triggers.

This work has two objectives. The first one is to investigate the capabilities of the open-source tool Cortex. By doing so, we aim at listing Cortex's medium and high-level capabilities. The second objective is to demonstrate the use of Cortex, and this is to be done thanks to the creation of a test environment for testing the capabilities of the tool. Meeting these objectives will allow the DYNAMO project to ensure that they are aware of all the capabilities of the tool, itself and in conjunction with the open-source tools TheHive and MISP and to ensure that the project makes the best use possible of these capabilities. When doing so, they will be able to make sure that all expectations are met in the development of the module of the task 4.2 when it comes to the use of Cortex.

Before going deeper into explaining what Cortex is and how it works, it is necessary to understand the choice MISP, TheHive and Cortex for the DYNAMO project. MISP stands for Malware Information Sharing Platform. It is a powerful open-source platform which can be seen as a community managed database storing threat intelligence. It essentially centralises information about threats called indicators of compromise (IOCs) which indicate that an observable is compromised. MISP allows its users to import new threat information and export existing threat information coming from MISP. This is where TheHive and Cortex come into play. TheHive is a Security Incident Response Platform which helps Security Operation Centre analysts (SOCs) managing security incidents and collaborating efficiently (TheHive Project 2020). The integration of MISP with TheHive permits the export to MISP of observables for which an alert was raise in TheHive and also the import of MISP events into TheHive as alerts. TheHive detects security incidents in observables thanks to the use of analysers which are external applications able to detect threats. These analysers can be accessed by integrating Cortex with TheHive. Finally, inside Cortex, there is a special analyser called "MISP" which allows early detection of threats by matching the given observable with IOCs contained in MISP. This is how powerful the integration of these open-source tools is.

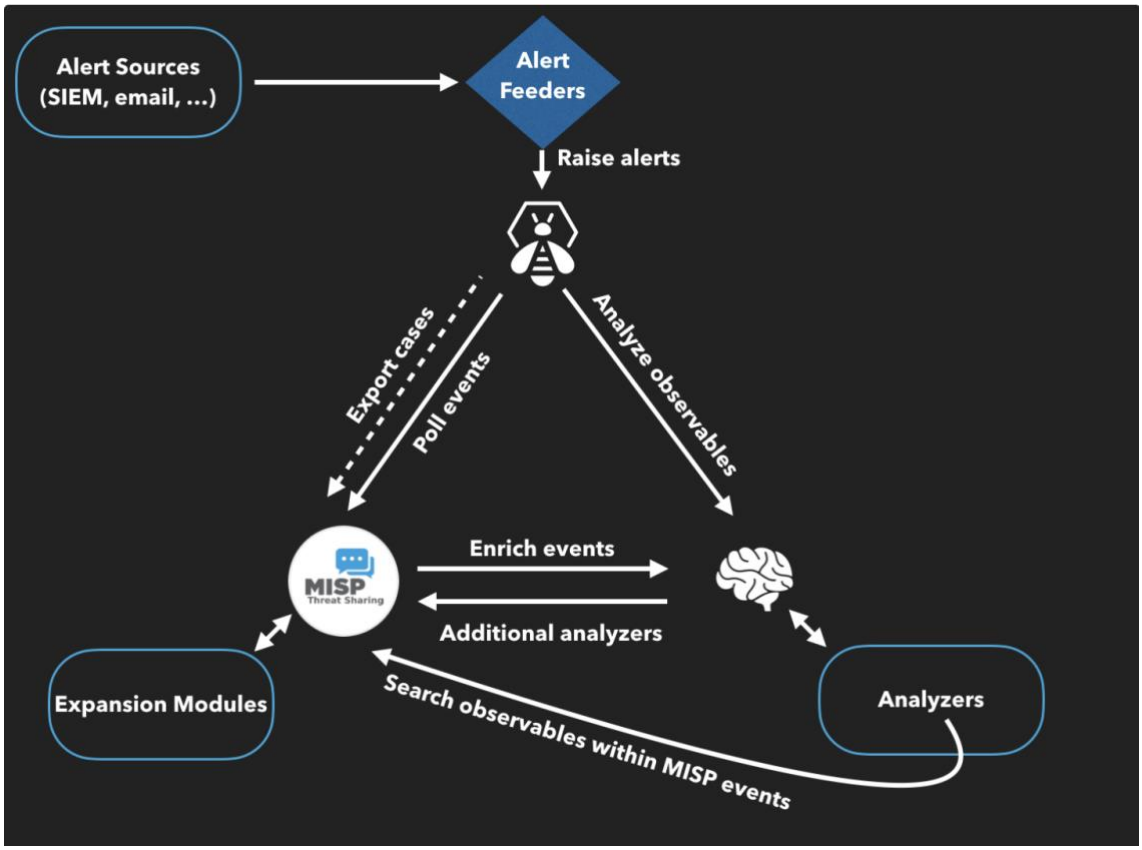

Figure 1 How Cortex, MISP and TheHive work together (Saâd Kadhi 2017)

The above figure summarises the way the three open-source tools work together.

In this document, we will first present the objectives, expected results, framework and methodologies used during this work. In the second part, we will look into the framework and draw requirements. The third part will be focused on describing the use of Cortex, an open-source tool used to analyse different types of observables to detect threats. We will start by describing Cortex, its purpose, and its capabilities. Then, we will analyse the integration of Cortex with the other open-source tools MISP and TheHive. This relates to the first thesis objective. Finally, we will demonstrate how Cortex can be used in practice. This part relates to the second thesis objective.

2 Background: Thesis objectives, expected results, framework, data collection and analysis methods

There are two objectives attached to this work. The first one is to detail and analyse all the capabilities of the open-source tool Cortex and the second one is to make a demonstration of these capabilities. These two objectives lead to two distinct expected results. The first expected result is a list of all the capabilities of the tool and an assessment of Cortex using a

theoretical framework and the second expected result is the creation of test environment testing these capabilities. In the report, we will present the different capabilities of the tool Cortex and will only be able to provide a step-by-step demonstration of the capabilities of the tool using illustrations.

In terms of theoretical framework, we will use the Computer Security and Incident Response Team (CSIRT) framework, compare the capabilities of the Cortex to the requirements of the framework and assess them. The CSIRT framework is particularly adapted because it focuses on incident handling which includes incident analysis. The purpose of Cortex is indeed to help analysts in the analysis phase of an incident handling.

To collect data for this work, we used ResearchGate, the ACM digital library and the open access journals of MPDI where we were able to find scientific articles supporting our analysis. We also used the documentation websites of the open-source tools Cortex, TheHive and MISP. We based our analysis on the CSIRT framework on two documents: the Handbook of Computer Security Incident Response Teams (CSIRTs) by MJ West-Brown and Computer Security Incident Handling Guide from NIST. Finally, we also collected data thanks to our own experience of using the tool.

Regarding the analysis method of the collected data, it was mostly qualitative. The collected data helped us better understand the capabilities of the tools studied and describe them in an orderly manner. We also assessed the capabilities of Cortex thanks to the CSIRT framework and provided grades based on qualitative information.

3 The CSIRT theoretical framework

In order to be able to analyse the capabilities and the use of Cortex, we need to compare it to widely used frameworks. This section aims at analysing the CSIRT framework and its requirements. In Summary and evaluation of the functionalities of Cortex, we will compare these requirements with Cortex's capabilities. In Results and outcomes (2/2): Setting up a test environment for Cortex, we will assess Cortex ability using use-cases based on requirements from the CSIRT framework.

The CSIRT framework aims at providing a guidance on the place of the CSIRT and its mission within its environment. Various services can be provided by a CSIRT team, however, there is one service that it must necessarily provide which is incident handling. (West-Brown et al. 2003). This section aims at describing the incident handling service.

The CSIRT can provide a variety of services that can be classified into "reactive services", "proactive services" and "security quality management services". "Reactive services" are a

category of services that comes as a reaction of requests, incidents, or threats. “Proactive services” come before any problem and intend to prevent them. Finally, “security quality management services” are services that are usually performed by other departments and where the CSIRT can provide a different viewpoint to the one of the traditional departments running the functions. However, to be called CSIRT, the team should at least provide an incident handling service. The incident handling service is part of the “reactive services” provided by the team. (West-Brown et al. 2003).

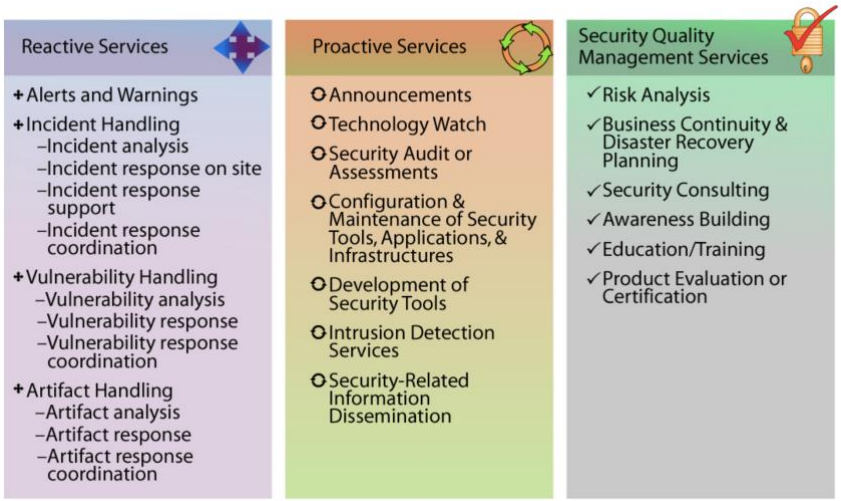

Figure 2 List of services that CSIRT can provide (West-Brown et al. 2003,25)

The Handbook of CSIRT provides the following definition of the incident handling services: “Incident handling involves receiving, triaging, and responding to requests and reports, and analyzing incidents and events.” The receiving and triaging activities along with other activities form part of the support activities which are essential activities that helps the CSIRT in the response and analysis part. In contrast, the analysis and response activities form part of the core activities of the CSIRT (West-Brown et al. 2003.).

3.1 Central activities

The Handbook of CSIRT divides the incident handling services provided by the CSIRT into four different services: “Incident analysis”, “Incident response on site”, “Incident response support” and “Incident response coordination” (West-Brown et al. 2003.).

Incident analysis is an essential part of the activities of the CSIRT. It is an activity that is necessarily part of the CSIRT programme which can contain one of more of the activities listed above. Its tasks can be separated into two levels. On the lower lever, based on the analysis of the information collected on an event or an incident, the team must determine its scope, the impact, provide a description of the incident and strategies for responding to it.

Results from the artifacts (1) and vulnerability (2) analysis, which can also be part of the CSIRTs services (see figure above), can be used in the incident analysis to help the CSIRT's team members have a better understanding of the incident and therefore provide a better analysis of it. On a higher level, the team's incident analysis service also includes a global analysis of all the incidents to understand trends and correlations. The incident analysis can include the collection of forensic evidence and also the tracing of an intrusion (West-Brown et al. 2003.).

As part of the incident handling services, CSIRTs might also be involved in incident response at the direct contact of users that suffered from an incident by analysing the impacted systems and repair them and recover information. The CSIRTs can also provide incident response support which would be a distant support as opposed to the on-site support described earlier. Finally, the CSIRT can coordinate the incident response between different stakeholders involved (West-Brown et al. 2003.).

The handling function of the handling of incident services provided by the CSIRTs has three main activities which are receiving incident reports, analysing these reports and the underlying information and determining responses and finally sharing these responses.

Requirements related to the handling function of the CSIRT are the following:

<u>In terms of incident life cycle handling</u>

- (3) Notification of impacted stakeholders when closing an incident ideally as part of the conversation ongoing during the incident
- (4) Reusage of the same tracking number when an incident is reopened and a new one if separate matter
- (5) Possibility to mark incidents as related and possibility to merge them
- (6) Ability to keep track of information related to the incident

<u>In terms of incident analysis</u>

There are two types of incident analysis. The "intra-incident analysis" analyses a particular incident while an "extra-incident analysis" makes an analysis of relationships between incidents. The ability to analyse incidents overall trends and to extract statistics is important for the CSIRT function. Getting, analysing and sharing this overall view of incident is necessary (West-Brown et al. 2003).

As for "intra-incident analysis", an analysis of the logs can be undertaken, and the requirements related to this are:

- (7) the ability to categorise them using multiple categories before the team receives it,
- (8) receiving the logs in a manner that it appropriate to its categories,
- (9) having the ability to authenticate the log,
- (10) cleaning the logs from sensitive information not necessary for the analysis
- (11) the ability to send extract of the log as incident follow-up to the relevant audience (West-Brown et al. 2003.).

Further requirements are provided by NIST and suggest to:

- (12) perform some profiling for networks and systems and recognize a normal behaviour
- (13) Create a policy that determines how long logs should be kept
- (14) Correlate events between one another and synchronize host clocks
- (15) Keep a "knowledge base of information"
- (16) Using search engines on the internet to research information about potentially abnormal activity
- (17) Use packet sniffers to gather more data
- (18) Have the ability to filter data by category of indicators
- (19) Request help from others (Cichonski et al. 2012.).

3.2 Support activities

These support activities are generally composed of the triage, feedback and announcement set of tasks or functions. The handling function has been described in the core activities section. We will further tackle the triage and the announcement functions as part of this thesis. The feedback function relates to the necessity to provide feedbacks to non-incident related requests (West-Brown et al. 2003.).

The below figure provides a view on the different functions.

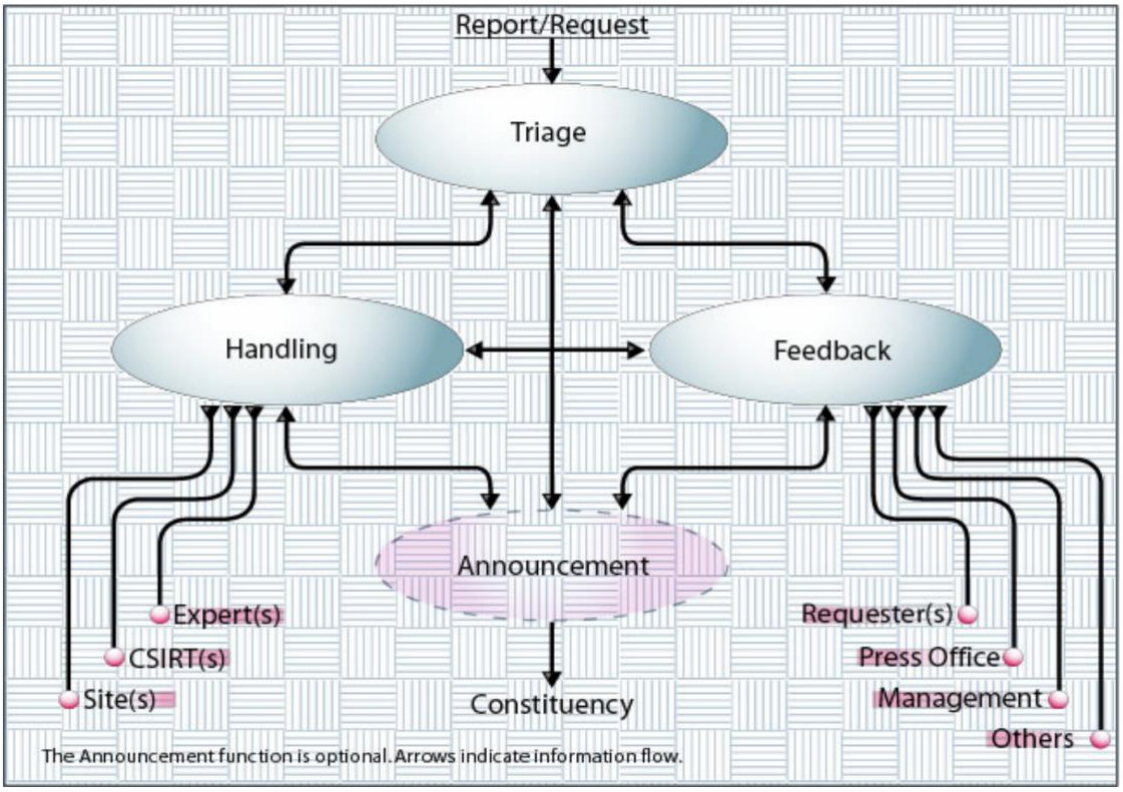

Figure 3 Different functions of the Incident Handling service (West-Brown et al. 2003,67)

- "Triage"

  As the point of entry, the triage centralises all the incident information coming from different sources. It then analyses it to see if it linked to past events and a level of priority is given. The Handbook of CSIRT defined different requirements for this function:

  - (20) Assigning and using tracking numbers
  - (21) Registering the contact information

- "Announcement" (22, 23, 24, 25, 26)

  This function consists in publishing to the dedicated public information about threats, providing details about an incident, informing about required measures to be taken to address a particular threat. We can distinguish between different types of announcements: "heads-up", "alerts", "advisory", "For Your Information", "guidelines" and "technical procedures". Several elements should be taken into account as part of the announcement function. The first one relates to the triggering of announcement, elements that trigger announcement should be clearly defined. Then, a categorisation may be done, that helps understanding which kind of announcement will be needed. Prioritisation of announcements should also be part of

the CSIRTs considerations, in order to ensure that the message has the right impact on its audience. The published information should be cleared and the channels for distributions should also be considered depending on the audience and the sensitivity of the information.

4 Results and outcomes (1/2): Understanding Cortex and its capabilities

This section aims at presenting Cortex and its capabilities. The list of the medium and high-level capabilities of Cortex can be found in Appendix1.

4.1 Cortex

4.1.1 What is Cortex?

Before we define what Cortex is, we need to understand what an open-source tool is. Open-source software (OSS) is code that is meant to be available to the public. This means that anyone can use it, study it, and modify it to what they want to use it for. The software is under a license that gives users the right and permission to use it for any purpose they want. This permission is given by the copyright owner. Hence, Cortex is an open-source tool which uses the AGPL licence (Affero General Public License). It is designed to analyse observables and generate active responses. It can handle different observable types such as email addresses, files hashes, URLs, and IP addresses and let the user choose among a variety of tools to analyse its observations. This is what really make the tool powerful. Cortex also provides the user with the results of these analyses (TheHive Project 2020).

In practice, the analysis of observables is done using external applications accessed by Cortex thanks the use of their APIs. These external tools are called “analysers”. Among these tools, we can find VirusTotal, Nessus, Censys.io and FortiGuard. In addition to this, Cortex also has a module for responders. Responders are actions that can be performed as a result of an analysis. Among the responders available, we can find Redmine, Palo alto minemeld, Umbrella blacklister and VirusTotalDownloader. While some of the analysers and responders are free to use and others require a license or a subscription (TheHive Project 2020). Cortex has 185 analysers and 50 available responders[1].

Automation of the actions performed by Cortex as well as the ability to send a lot of observables at once is possible via the use Cortex4py which is an API client for Cortex (TheHive-Project 2021).

[1] At the time of writing this paper.

4.1.2 Cortex analysers

An analyser helps analysing observables such as hashes, IP addresses, files, domain names and URL. As all these analysers can be found within one tool, this helps automating a big part of the tasks of SOC analysts which is the search for tools to analyse observables (TheHive Project 2020). Among these analysers we have freely accessible analysers such as Abuse finder, Cybercrime-Tracker and analysers that are license and subscription based such as Nessus and VirusTotal.

- VirusTotal Scan analyser

VirusTotal analyser can check for suspicious files, URLs, domains. When a malicious content is found, it is automatically shared with the VirusTotal community. The files, domains, URLs shared with VirusTotal are inspected and scanned with more than seventy antivirus scanners such as AVAST Software Avast Antivirus, Fortinet, F-secure. VirusTotal also has antivirus website domain scanning tools such as Alienvault cybercrime, and Comodo Site Inspector.

- urlscan.io Search analyser

urlscan.io is a free public website that helps analysing all possible HTTP connections, site content, domain tree, and the relation with other sites. Urlscan.io gives the ability to see the backend of a website. It opens a sandbox in the background and browses the website that the user has provided. It provides the users with all the redirected links to the website, and it also detects the technologies used by the website. It also helps knowing how many times a specific website has been scanned on urlscan.io. urlscan.io can analyse different observables such as Ips, hashes, and domains.

- Phishing tank analyser

Phishing tank analyser is a free website where anyone can verify, track, share and submit phishing data. Cisco Talos Intelligence Group (Talos) runs this site with the purpose of sharing the information or phishing data freely with everyone who may it. Phishing tank analyser makes an API call to the Phishing tank database to query information on URLs that have been identified as phishing website. For this analyser to be effective, the user must sign up to the Phishing tank website to get an API key that will allow him/her to make API request to Phishing tank database (Phishtank, no date).

Cortex also encourages users to create their own analysers. To do so it is necessary to understand what analysers are, from a technical standpoint. Analysers are program able to take observables along with information related to the configuration of the analyser as inputs

and to give an output to the user after having performed an analysis of the given input. The input and output are in a JSON format.

The Cortex documentation specifies the necessary elements for the creation of an analyser. The creator needs to provide:

- The program which could be written in any language as long as it is supported by Linux (if the program is in Python, then a requirement file with dependencies for the program is required)
- Information about the configuration of the program including the observable types it is able to analyse, the creator of the analyser, Traffic Light Protocol[2] and Permissible Action Protocol[3].

4.1.3 Cortex responders

A responder provides actions that are performed as a result of an investigation or an alert. This response can be to block an IP address that is determined to be malicious. Analysers are meant to collect information while the responders are meant to provide an action as a result of the information collected. Cortex has free responders such as Redmine, Wazuh, and Palo Alto Minemeld. Cortex also have responders that are based on subscriptions and licences such Umbrella blacklister, VirusTotalDownloader and KnowBe4 (TheHive Project 2020).

- Redmine responder

Redmine responder is used in Cortex to create a ticket in Redmine ticketing system from a case occurred in Cortex. To create a ticket, Redmine ticketing system will use the title, subject, and the case description from Cortex as body of an issue.

- Palo Alto Minemeld responder

The Palo Alto Minemeld responder is a responder that sends observables from Cortex to Palo Alto Minemeld. Palo Alto minemeld is an open-source application from Paloalto network that allows users to share threat intelligence (Nils Kuhnert 2020).

- VirusTotalDownloader responder

VirusTotalDownloader is a responder that allows users to download on VirusTotal website a sample of a malware when a user submits a hash. For this responder to work, the user needs

[2] Traffic Light Protocol are described in the Cortex demo section
[3] Permissible Actions Protocol are described in the Cortex demo section

to have a premium valid API key that he/she will get from the VirusTotal website (TheHive Project 2020).

It is also possible to create responders. The creation of a responders resembles the creation of an analyser. Responders are program able to take a JSON input and to give an output to the user after having performed an action with the given input. The input and output are in a JSON format.

The Cortex documentation specifies the necessary elements for the creation of an analyser. The creator needs to provide:

- The program which could be written in any language as long as it is supported by Linux (if the program is in Python, then a requirement file with dependencies for the program is required)
- Information about the configuration of the program including the types of data accepted by the responder, the creator of the analyser, Traffic Light Protocol[4] and Permissible Action Protocol[5].

4.2 Integrating Cortex with TheHive

4.2.1 What is TheHive?

TheHive is an open source "Security Incident Response" platform which helps SOC analysts and other security analysts manage security incidents and threats. It helps professionals collaborate by allowing them to work on cases and communicate on the resolution of security issues. TheHive has the ability to integrate with other tools such as MISP and Cortex (TheHive Project 2020).

When we successfully manage to install TheHive and before we start using it, we must create an administrator account. This administrator account will allow us to create an organisation, i.e., a group of users with different privileges or user right accesses and assign users to this organisation. We can assign different roles to a user. Amongst these roles, we have a platform administrator, an organisation administrator, or an analyst. These roles give users a different type of access to TheHive's functionalities.

The functionalities of TheHive can be described by imagining that an alert comes into the organisation's network. An alert is a notice that there might be a danger and that it might require an action. The alerts in TheHive can be generated from another platform and through

[4] Traffic Light Protocol are described in the Cortex demo
[5] Permissible Actions Protocol are described in the Cortex demo

APIs. To create an alert in TheHive, users need to have a manage alert permission. Alerts can contain observables. Based on these alerts, users can create cases in TheHive for investigation or sharing purposes.

In TheHive, a case is an issue that may need an action or an investigation. It can be created from scratch or using a template. TheHive's cases can also be merge with one another to make one single case. A new case contains the following details: a title, a severity level, a date, tags, a description, amongst others.

A case can be broken down into different sub-cases called tasks. These tasks can be assigned to other users. TheHive provides a user with the ability to see the list of tasks assigned to him/her. He/she can then check the details all the available tasks and start working on one of them. TheHive indicates whether a task has been started or completed.

A case can include observables. These observables can be of different types such as IP, file, filename, hash, and domain. When creating an observable, TheHive also requests the user to include tags and a description. After having created the observable, we can see if it is related to other cases that we have been working on before and the details related to it.

Users are thus able to create cases, assign a task to another user and run analysers. These functionalities help users collaborate with each other and share information efficiently within an organisation. The collaboration in TheHive happens in the real-time, this means that users can see when a task has been created and who created it thanks to TheHive live feeds. TheHive also provide users with a search option that allows them to search in TheHive for cases, tasks, tasks logs, observables, and alerts.

Anand and Joris (2021) in their analysis, attributed a rating of 3 out of 5 to TheHive and concluded that TheHive could be considered a mature Security Incident Response platform.

4.2.2 Why is the integration of Cortex with TheHive beneficial?

We cannot talk about TheHive without mentioning the results we could get by integrating it with Cortex. TheHive allows users to make API calls to Cortex to use analysers that are available in Cortex to investigate or analyse observables. This is possible without leaving TheHive. Indeed, TheHive has the ability to let the user choose the analysers available in Cortex, to let him/her run the analysis and to retrieve the results from the analysis and to display them in TheHive (TheHive Project 2020). This is key because it then allows the user in TheHive to make decisions based on the results of the analysis and to communicate them to the other members of his/her organisation.

In addition to that, TheHive has a special module allowing it to make the results of the analysis from the Cortex analysers more readable. Indeed, the results of the analysis are

initially in a raw format which can be hard to quickly analyse and therefore to make decisions based on them. TheHive also solved this problem, making the communication inside the organisation about incidents seamless.

Finally, TheHive is also able to make use of another Cortex module which is the Cortex responders. In TheHive, it is possible to create automations which allow the users to respond to incidents. This is done my using the available responders in Cortex. This is what makes Cortex an active response tool.

4.3 Integrating Cortex with MISP

4.3.1 What is MISP?

MISP is an open-source software that essentially gathers malware information. It can be seen as a database containing important information called IOCs. This database is community-based, i.e., available to the MISP community which has the possibility to enrich it by importing new IOCs (depending on sharing levels). MISP can also interface with other tools thanks to its API and allow the tools to access its database. By accessing the MISP database, the tools can analyse new observables and potentially find matching events related to these observables in the MISP database (TheHive Project 2020). We will briefly describe a central element of MISP which is the events.

Inside MISP, once we have created a user, we have the possibility to enable feeds to have access to events that the MISP community has shared and that are extracted and gathered from different sources such a blog posts or other external organisations. An event is information that is shared about a malware, ransomware and malicious groups which contains metadata around these events. This metadata is of different types can be summarised as:

- the creator organisation, i.e., the organisation that initially collected information around the event,
- the tags associated with the event which helps classifying events,
- the level of the threat,
- the status of the analysis,
- a description of the event,
- attributes related to the event which are general information related it, and,
- IOCs or observables related to the events.

In addition, a great feature of MISP is the ability to access related events. Indeed, MISP has the ability to link different events.

Two features that make MISP a reliable collaboration platform for the DYNAMO community are its sharing model and its “proposal” feature. MISP allows users to define the sharing level

of events they share. Four levels are indeed available: sharing only with own organisation, sharing within the community, sharing with connected communities, or sharing all the communities in MISP. We can imagine that different companies from critical sectors inside the DYNAMO community would create their own organisations and be part of the same DYNAMO community within which they would share most of the events but could also collaborate with the broader MISP community. MISP collaboration also relies on another feature which is called “proposals”. Proposals ensure that data in MISP is truthful and complete by allowing users to make suggestions on events that are created by other organisations. The creator of the event can then choose to accept or not accept the proposal. If the proposal is accepted, the event is then updated (Iklody et al. 2016).

Davy and Wouter (2021) show that to reduce the potential for false positives and negatives in the cyberthreat detection by the security analysts, the sharing of IoCs could be further improved by sharing encrypted machine learning models and related artifacts as IoCs used to spot cyber security attacks more effectively, in a more precise and detailed way.

4.3.2 Why is the integration of Cortex with MISP beneficial?

Integrating MISP with Cortex consists in setting up the MISP analyser in Cortex. This MISP analyser allows users of Cortex to compare their observables with the rich database of IOCs maintained by the MISP community and detect threats within the analysed observables. The MISP analyser is able to analyse a lot of observable types such as domain names, IP addresses, URLs, hashes, mail and others and retrieve the MISP events associated with the observables. An integration can also allow Cortex’s analysers to be accessed in MISP (TheHive Project 2020).

4.3.3 Further thoughts on the integration of MISP and TheHive

As discussed in the introduction, a direct integration of TheHive and MISP is also possible. It goes beyond the scope of this paper, but we can briefly describe what it allows. This integration makes an important feature of MISP available, which is importing new threat intelligence. Indeed, the link between TheHive and MISP allows analysts to share observables for which an alert has been raised and to let the MISP community know about the existing threat. Members of the MISP community will then be able to detect threats earlier. The integration would also allow MISP events to be imported into TheHive. These MISP events will be seen as alerts in TheHive with their observables.

5 Summary and evaluation of the functionalities of Cortex

The list of requirements from the CSIRT framework and the assessment of Cortex according to the CSIRT framework was detailed in Appendix 2.

From this assessment, we note that Cortex is able to meet a restricted number of requirements from the CSIRT framework (15.4%). These requirements relate to the "artifact analysis" and "incident analysis" functions of the CSIRT "incident handling services" and concern:

1. Using results artifacts analysis (1)
2. The need to keep a "knowledge base of information" (15);
3. The need to use search engines on the internet to research information about potentially abnormal activity (16) and;
4. The need to request help from others (19).

We have evaluated that the ability of Cortex to meet these requirements was between 3/5 and 5/5 with an average of 3.75/5 or 75%.

This demonstrates that the capabilities of Cortex are very specific. Indeed, Cortex's abilities only revolve around the analysis of observables (or artifacts) and the generation of active responses.

6 Results and outcomes (2/2): Setting up a test environment for Cortex

Now that we have describe the capabilities and uses of Cortex and its integration capabilities with the other open-source tools which are part of the DYNAMO project, we will spend some time on a demonstration of Cortex.

Due to technical problems linked to version incompatibilities of some tools required for the functioning of MISP and Cortex, we will only provide a demonstration of Cortex itself and its use in conjunction with TheHive.

6.1 Working with Cortex

6.1.1 Creating a VirtualBox

VirtualBox was used to install and demonstrate the capabilities of the open-source tools Cortex and TheHive. VirtualBox is an Oracle virtual machine that helps users to create and to use different operating systems such as Linux, Windows, Mac OS on a virtual software platform (Oracle 2022). This allows users to run multiple different virtual machines on a

single physical machine without the need to modify the existing operating system that is installed on the physical device. For this demo, we used Ubuntu 20 LTS Linux distribution.

6.1.2 Installing Cortex

The programming language used to write Cortex is Scala. Angular JS was used for the frontend. The analysers contained in Cortex are in Python but to write an analyser any programming language that Linux supports can be used.

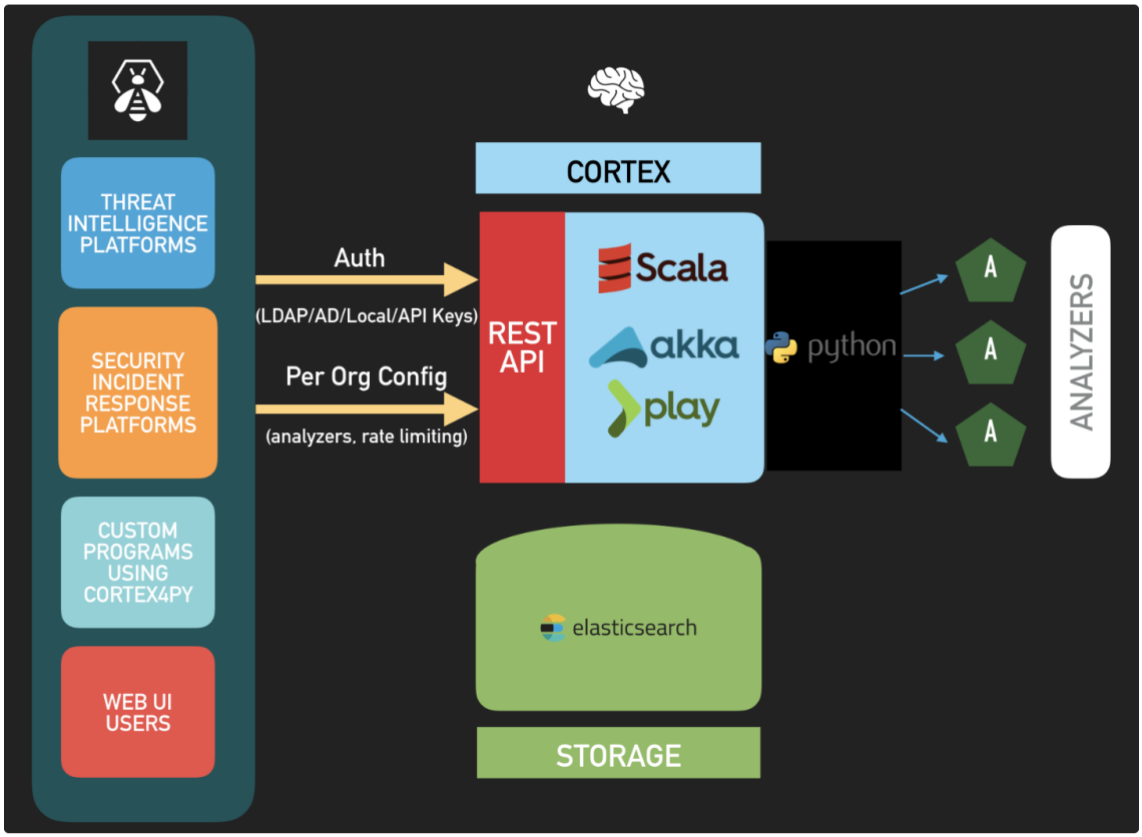

Figure 4 Architecture of Cortex (Saâd Kadhi 2019)

Once the VirtualBox is set up and we have created the user profile, we need to open the terminal to start installing Cortex. Before being able to install Cortex, we need to install Elasticsearch which is used as a database for Cortex. Elasticsearch listens on port 9200. After this, we can start installing Cortex. The installation of Cortex is done in two steps: installing Cortex itself and installing Cortex's analysers and responders. Cortex listens on port 9001.

For the demo, Cortex version 3.1.4-1 and ElasticSearch version 7.9.1 were used.

6.1.3 Cortex demo

The demonstration of the use of Cortex focuses on two uses-cases that are linked to the capacities of Cortex:

1. Separating roles within an organisation
2. Effectively managing analysers and responders

1. Separating roles within an organisation

To start with, in the first installation of Cortex, we are prompt to create a user. This user has super admin privileges and is able to create new users and organisations. Cortex comes with a default organisation and the member of this organisation cannot run or configure analysers because of the super administrator role. For this reason, we have to create a new organisation when we start using Cortex and, in this organisation, we will create users that have different roles. To create a new organisation in Cortex we click on the “Organisations” tab.

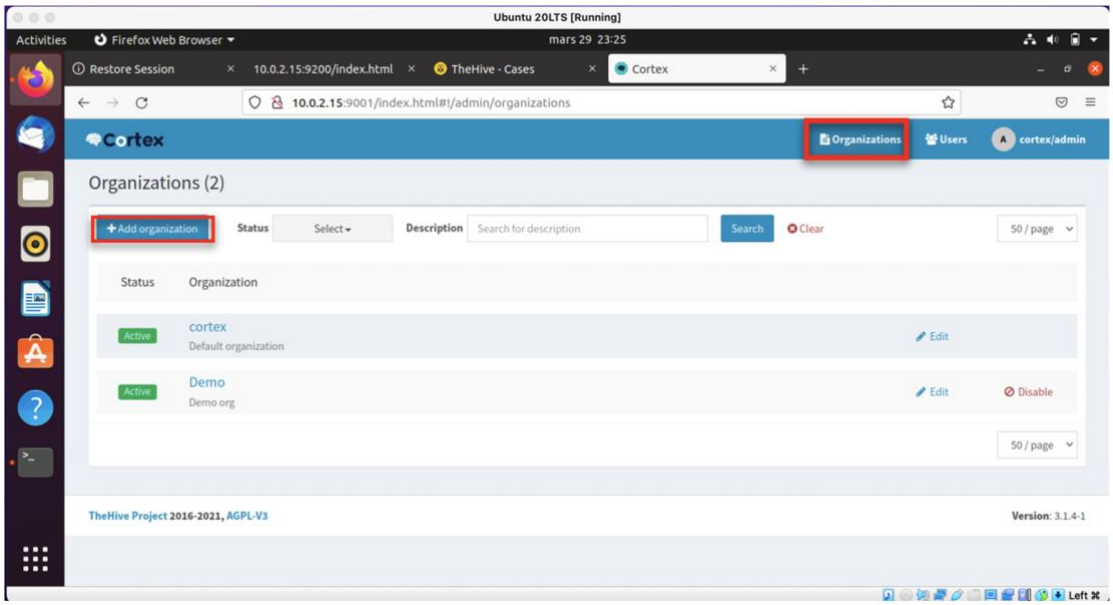

Figure 5 Cortex: adding an organisation

The administrator is the only one who can create organisations. We click on “Add organisation” to create a new organisation and after this, we can provide the organisation’s name and description.

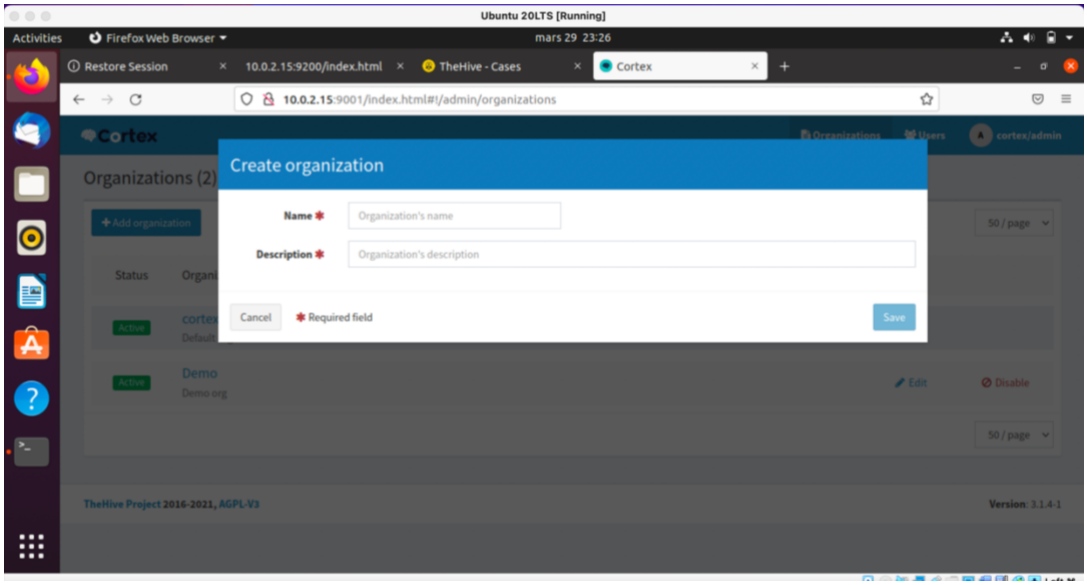

Figure 6 Cortex: providing information to create an organisation

We can create new users in the Cortex platform by clicking on the “Organisation” tab, and after this, clicking on “Add user”.

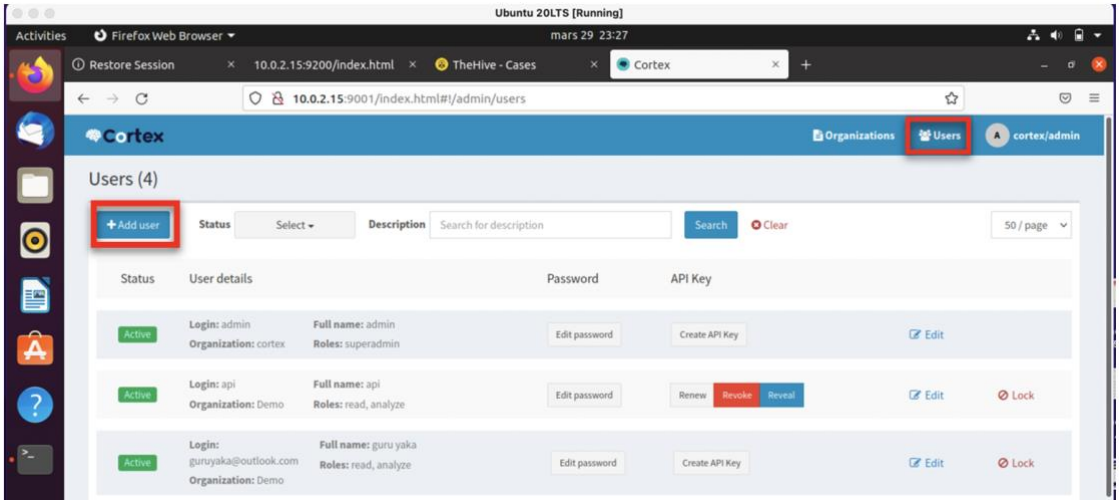

Figure 7 Cortex: adding a user

After this, we can provide the user’s login name, the user’s full name and its role. We have different roles to choose from that will determinate what a user can do. Among these roles we have:

- read only,
- read and analyse; and finally, and
- read, analyse, and being an organisation administrator.

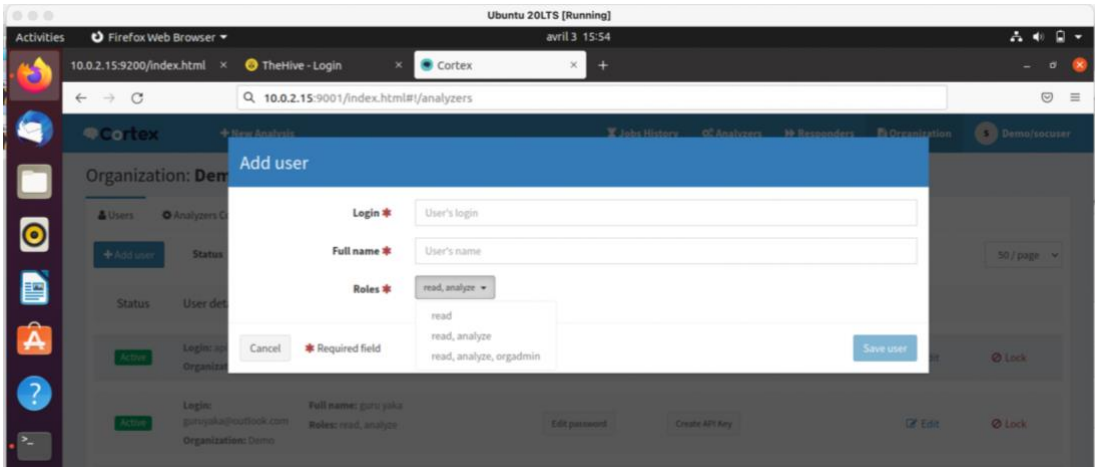

Figure 8 Cortex: selecting the user role

After creating a user, we can give it a password by clicking on the “Edit password” button.

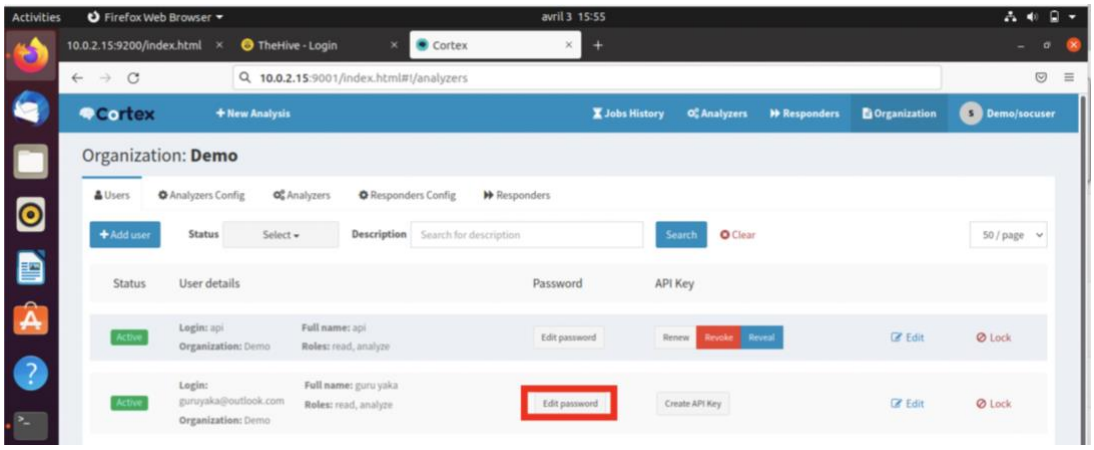

Figure 9 Cortex: providing a user with a password

2. Effectively managing analysers and responders

Cortex has 180 available analysers that users can enable. To enable analysers, we must click on the “Organisation” tab, and then on the “Analysers” tab. There, we will see a list of all the available analysers.

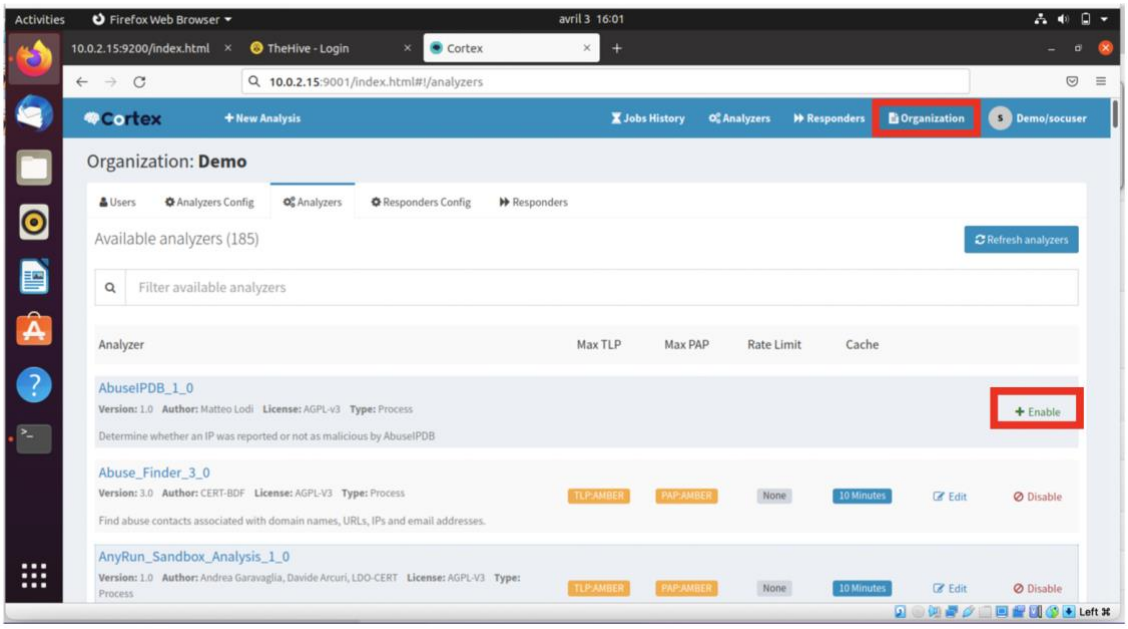

Figure 10 Cortex: seeing available analysers

From the list of the available analysers, we can choose which analysers we want to enable or disable. By default, all the analysers are disabled, and we need to enable them for them to work. For this demonstration, we choose an analyser named “AbuseIPDB_1_0”. To enable this analyser, we need to provide an API key. This API key is mandatory, and we have to get it from the AbuseIPDB. After providing an API key, we have other information to provide. The mandatory information is marked with a red star. We click on the “Save” button to save the information that we have provided.

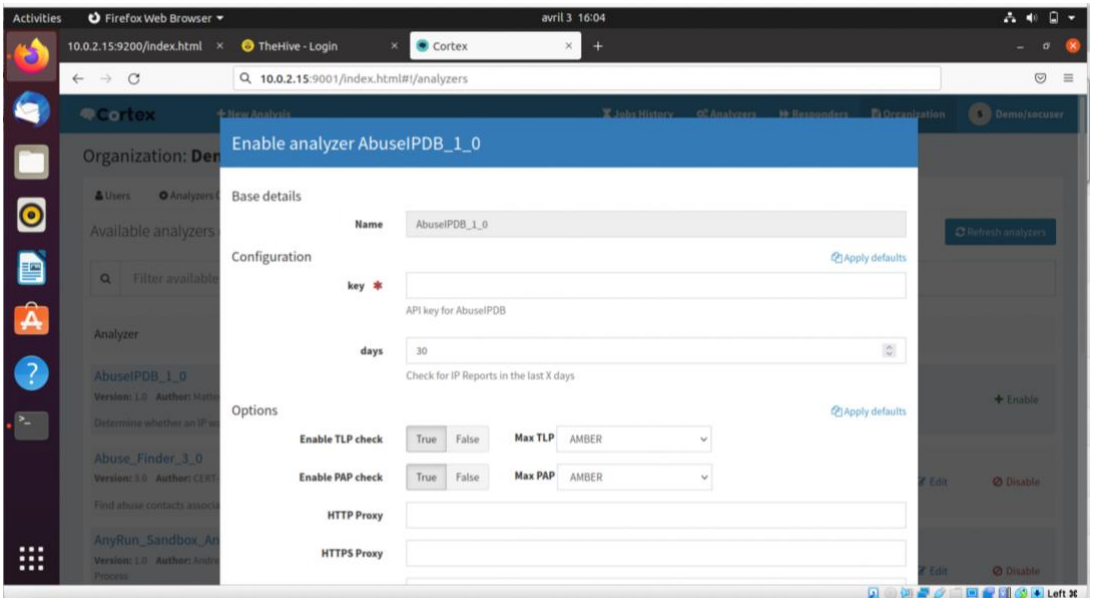

Figure 11 Cortex: enabling an analyser

Just like analysers, responders are, by default, disabled. We first need to enable them. To enable the responders, we click on the “Organisation” tab, and we click on “Responders”. A list of all the available responders will appear.

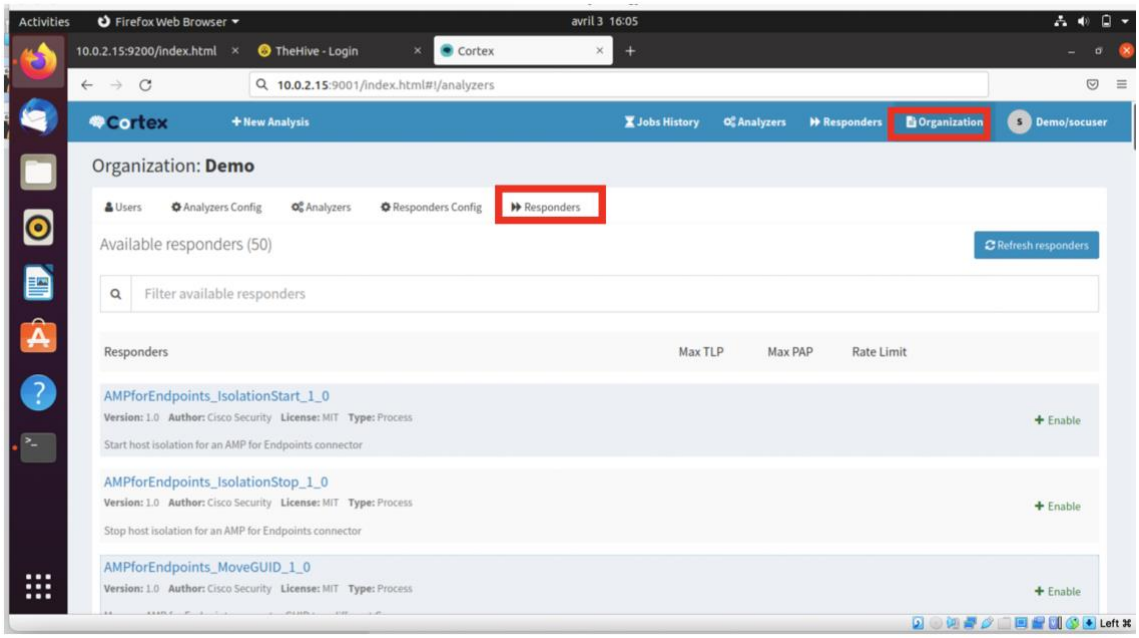

Figure 12 Cortex: seeing available responders

To enable a responder, we have to click on “Enable” and we have to provide the information requested such as the API key and other mandatory information specific to the responder. After this, we click on the “Save” button the save the information. This will enable the responder.

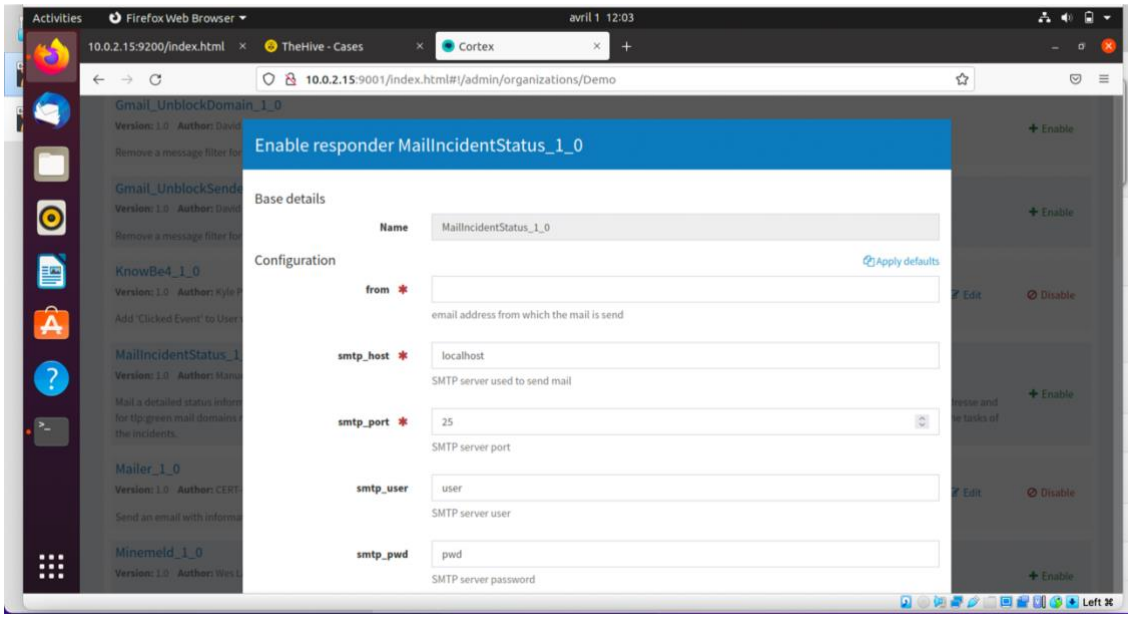

Figure 13 Cortex: enabling a responder

1. Running analysers and responders

When we are logged in as a user with “read and analyse” abilities, we can run analysis by clicking on the “New analysis” button on the upper left corner.

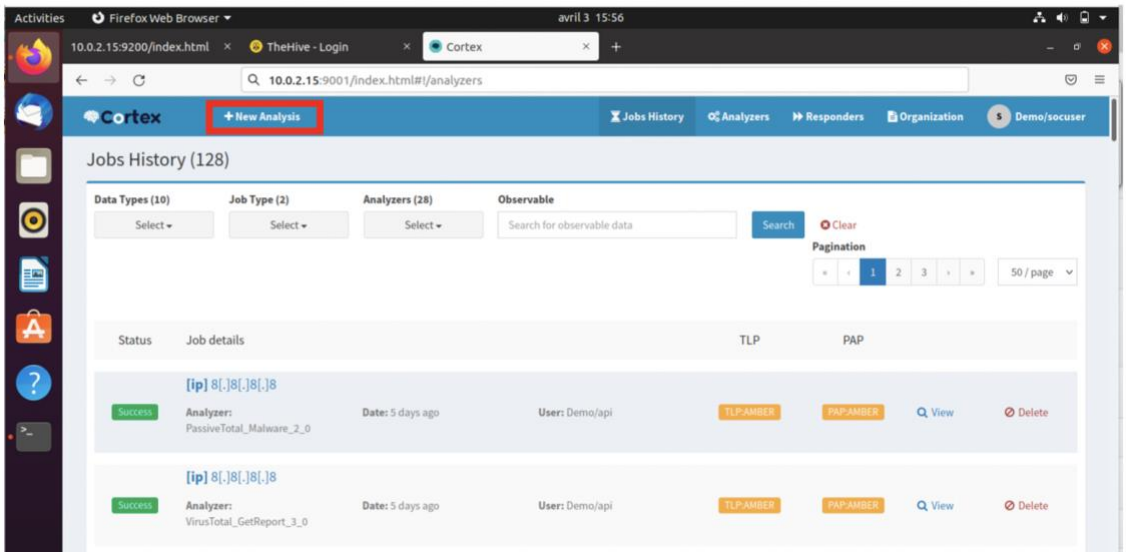

Figure 14 Cortex: creating a new analysis

After clicking on new analysis, a “Run analysis” window will prompt us to provide the following information:

- the TLP (Traffic Light Protocol) which indicates when the information can be shared and to which audience. It helps protecting sensitive information from being shared with the wrong audience.
    - Red is for “not for disclosure, restricted to participants only”
    - Amber is for “limited disclosure, restricted to participants’ organisations”
    - Green is for “limited disclosure, restricted to the community”
    - White is for “disclosure not limited”
- the PAP (Permissible Action Protocol) which indicates how the information can be used by the analyst receiving it.
    - Red is for “non-detectable actions only”
    - Amber is for “passive cross-checks”
    - Green is for “active actions allowed”
    - White is for “no restrictions in using this information”
- the data type which corresponds to the data type of the observable that we want to analyse. This data type can be an IP, domain, URL, and others, and,
- the observable itself.

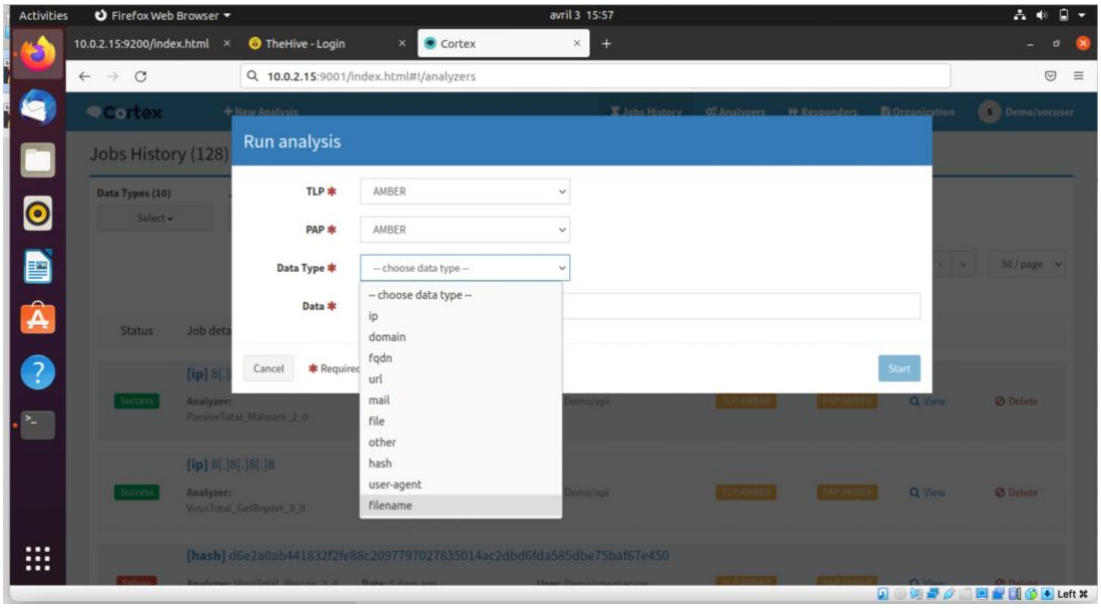

Figure 15 Cortex: providing analysis data (1/2)

Based on the data type, Cortex let us choose which analysers we want to use. Indeed, each analyser has observable types that it is able to analyse. After providing all the details, we can click on “Start” to proceed with the analysis.

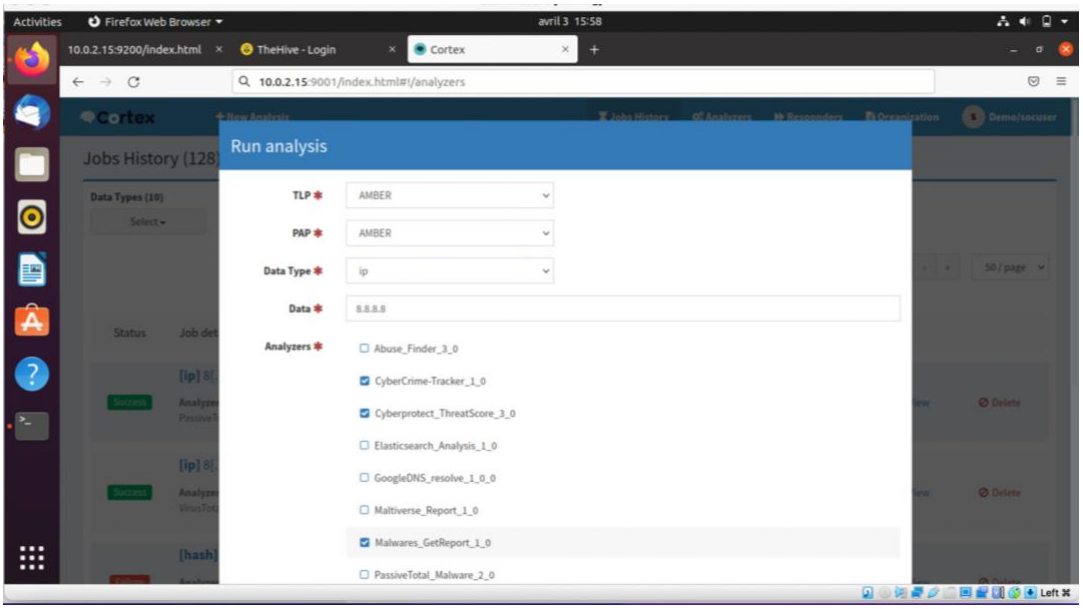

Figure 16 Cortex: providing analysis data (2/2)

We can also run the analysers by clicking on the “Analysers” tab and we choose an analyser to run an analysis with. The choice of the analyser will depend on the type of our observable.

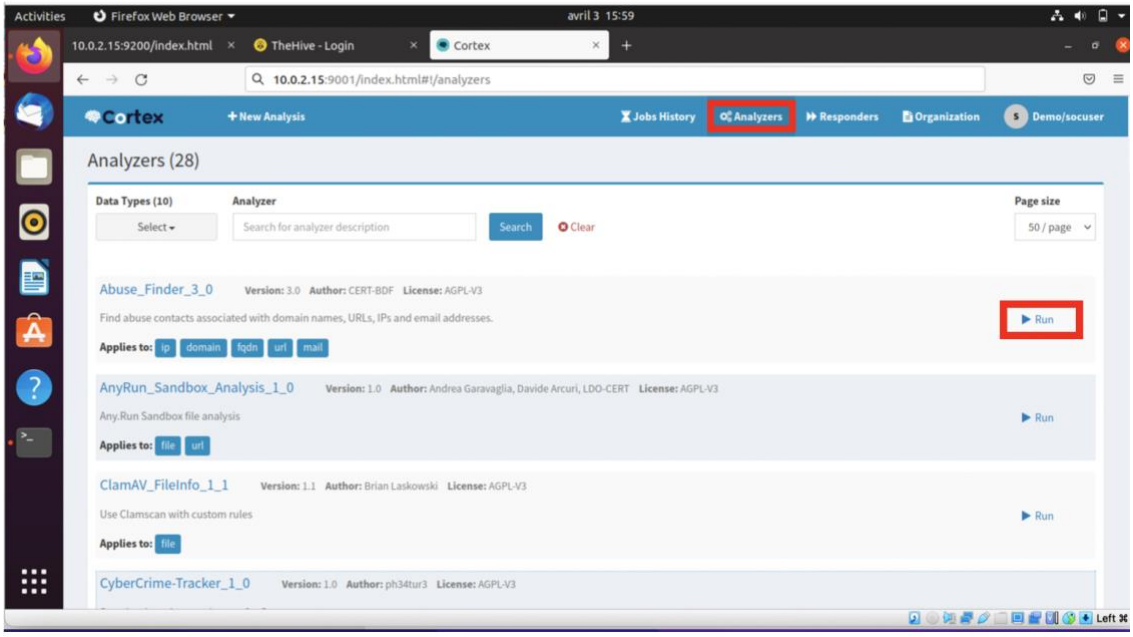

Figure 17 Cortex: running an analyser

After running the analysis, we can see its status. This status can be in “Progress”, “Failure”, or “Success”.

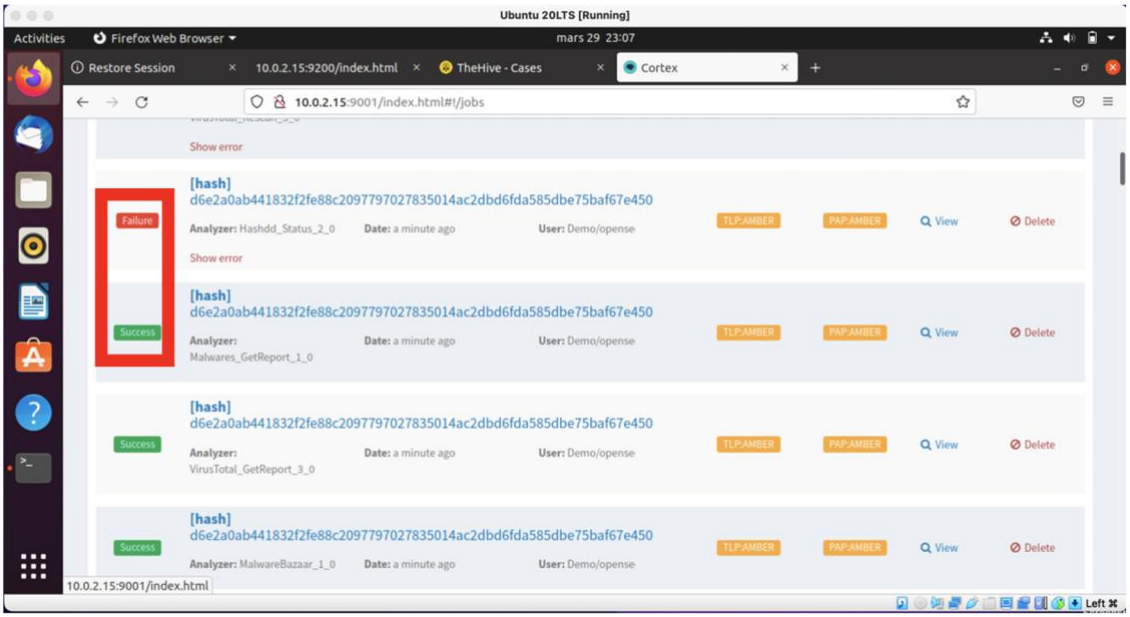

Figure 18 Cortex: checking the status of the analyses

When the analysis finishes, we can view its details and we also have the ability to delete the analysis if necessary. We click on “View” to see the job report. In the report, we can see the details of the analysis.

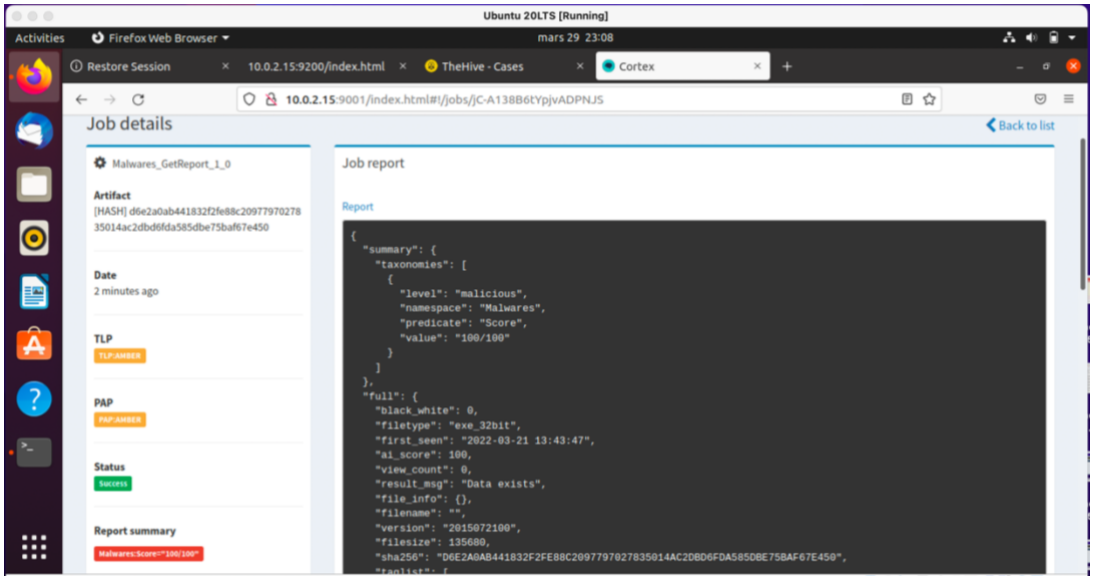

Figure 19 Cortex: seeing an analysis report

From these two use-cases, we can conclude that Cortex is effectively able to separating roles within an organisation thanks to its different roles. The tool also allows for an effective management of analysers and responders as demonstrated by the possibility to enable analysers and responders, to run analysers and to access analysis reports.

6.2 Working with TheHive

6.2.1 Setting up TheHive

Like in the case of Cortex, before we install TheHive, it is necessary to take an extra step and install Cassandra, which is a backend database for TheHive. In our installation, we also need to make sure that we have the correct version of Java. Cassandra listens on port 9042 and TheHive on port 9000.

For the demo, TheHive version 4.1.18-1 was used.

6.2.2 Integrating Cortex with TheHive

In order to integrate Cortex with TheHive, we first need to create an API key for our user in Cortex (Web UI) by clicking on the “Create API Key” as shown below.

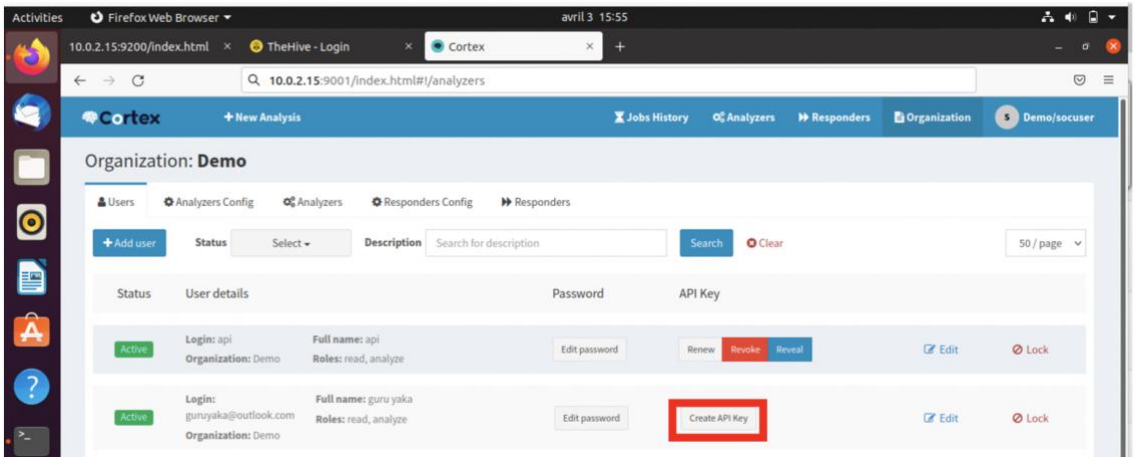

Figure 20 Cortex: creating an API key

Then, we need to include the API key in the configuration file of TheHive. This will allow the connectivity between the two tools. In TheHive (Web UI), we can go to “about” and we will be able to confirm that the connection was established.

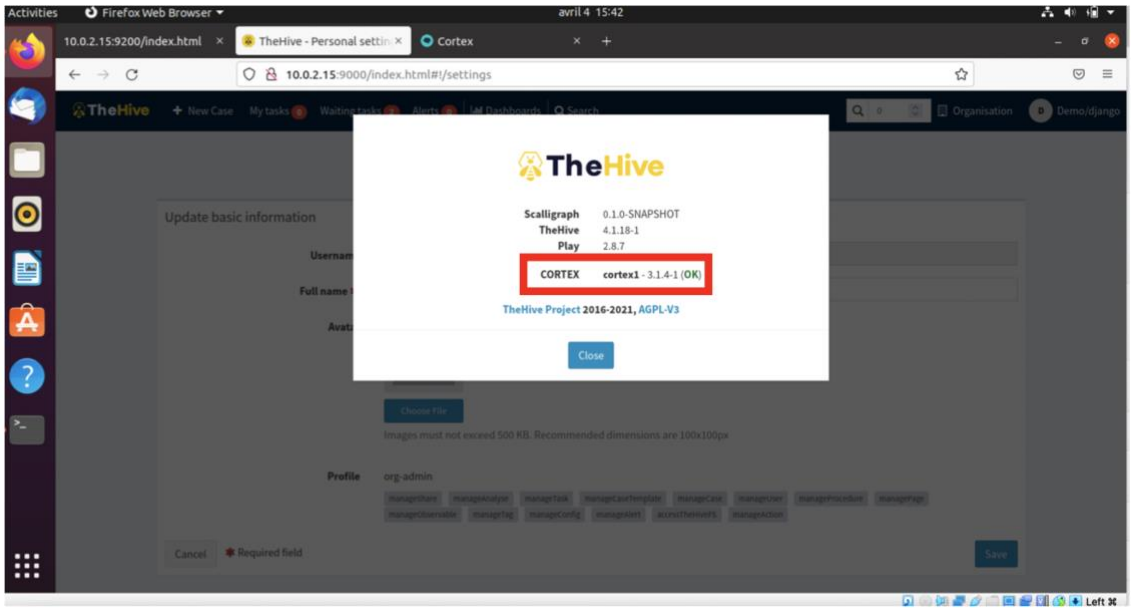

Figure 21 TheHive: checking the status of integration of Cortex with TheHive

6.2.3 TheHive demo

The demonstration of the use of TheHive focuses on two uses-cases that are linked to the capacities of the tool:

- Separating roles within an organisation
- Managing the “incident life cycle”

1. Separating roles within an organisation

When we successfully install TheHive, the first thing to do is to create an administrator account. This account will help us to log in to the TheHive. There are certain things that an administrator can do and that a normal user cannot do such as creating a new organisation. To create a new organisation in TheHive we have to expand the “Admin” tab, and we click on “Organisations”. This allows us now to click on “New Organisation” to create an organisation. In this demo, we have created an organisation called “admin”.

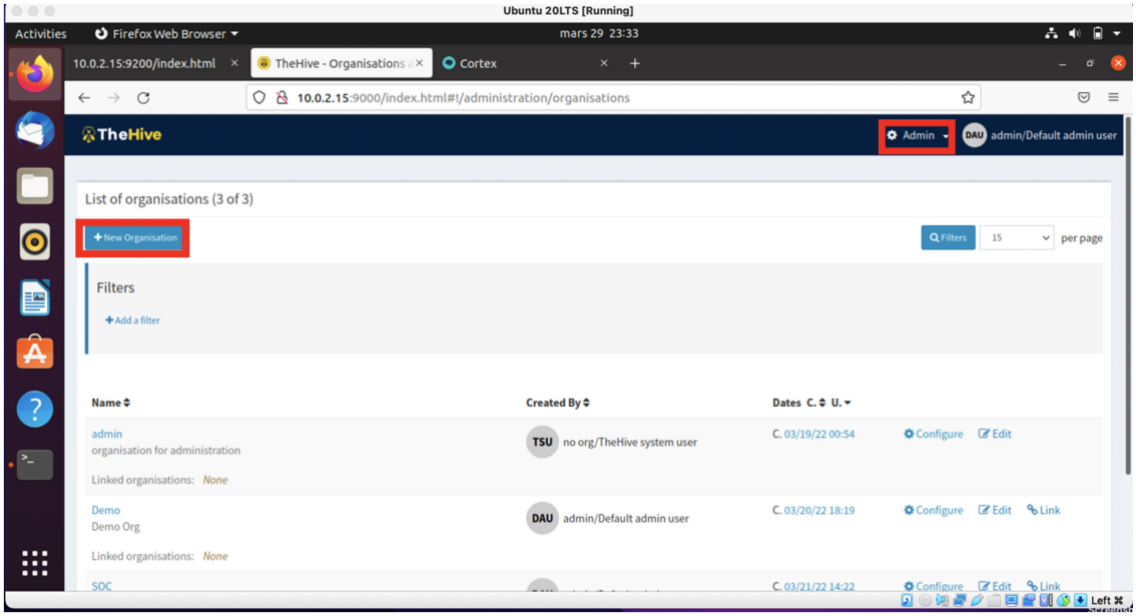

Figure 22 TheHive: creating a new organisation

We notice that a list of all the organisations is showing up on the screen with other organisations previously created such as “Demo” or “SOC”.

After having created an organisation, we can go on to create a new user for the organisation. To do so, we click on the organisation to which the new user will belong. In our case, we click on the “admin” organisation. We can see all the users belonging to the “admin” organisation in the “Users” tab. We can create a new user by clicking on “Create new user”.

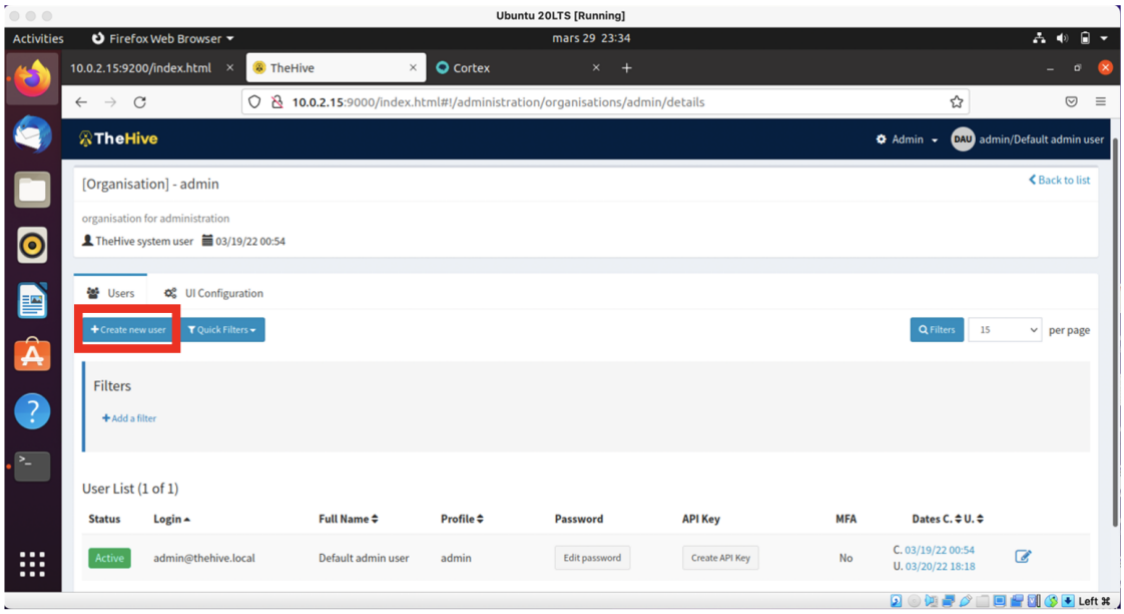

Figure 23 TheHive: creating a new user

When creating a user, we can choose the “profile” of this user. A profile in TheHive is category of permissions that will be assign to new user. This category will determine what a user can and cannot do. There are four built-in profiles which are: admin, analyst, organisation admin and read only. Each one of these profiles has specific permissions (also known as “user rights”). We can also create custom profiles and provide them with the permissions we want.

To create custom profile, we click on “Admin” tab and then “List of profiles”. We then click on the “New Profile” button. We can also edit or delete existing profiles on the list.

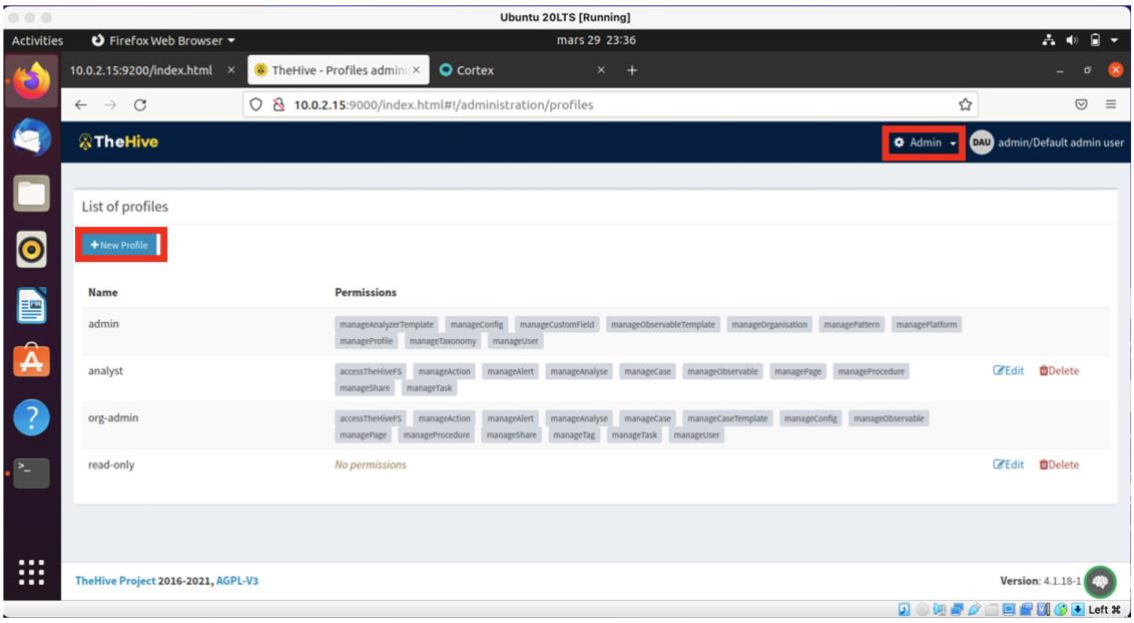

Figure 24 TheHive: viewing the list of profiles

In TheHive, everything is a case. A case is a situation that needs an action or an investigation. We can neither create cases nor run analysis when we are logged in as an administrator. To create cases and run analysis we need to be logged in as a user.

We are going to explore a case scenario where we noticed an attack from certain IP address on our mail server. In this scenario, we are going to go through the creation of a new case, the assignment of tasks to other users, the running of analysis using Cortex analysers and the closure of the case we opened.

2. Managing an “incident life cycle”
    - Creating a case

To proceed with this scenario, we start by creating a new case. To create a new case, we click on “New Case” on the upper left corner.

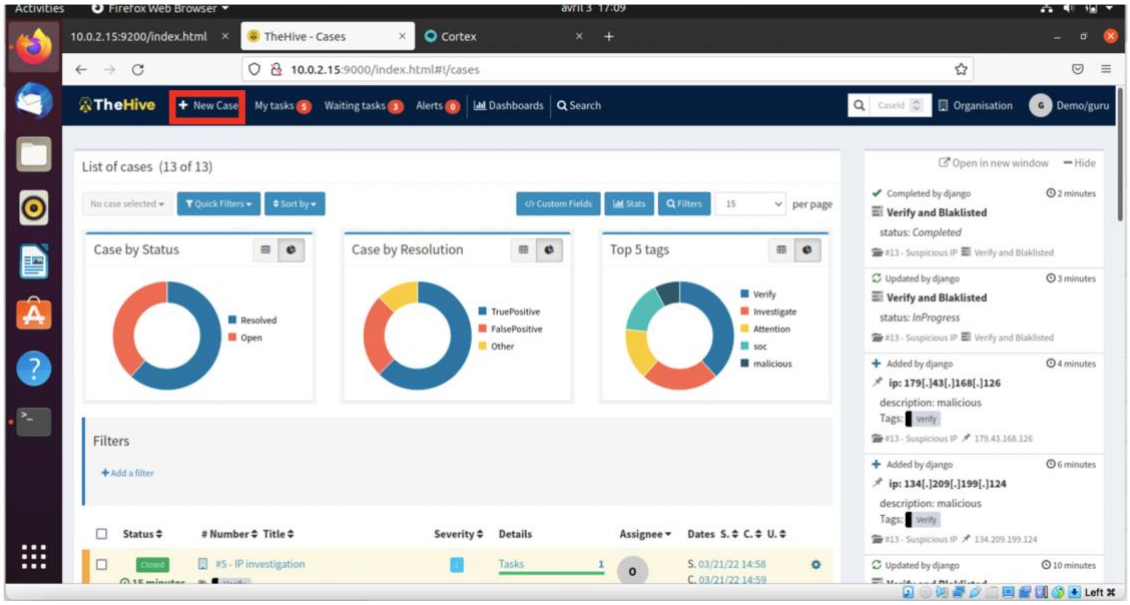

Figure 25 TheHive: creating a new case

A “Create new case” window will pop up where we have to provide the details about the case. The details we need to provide in order to create a case are:

- the title,
- the date which is automatically assigned to the date of case creation but can changed to the date we want,
- the severity of the case (L= Low, M=Medium, H=High, C=Critical),
- the TLP,
- the PAP,
- the tags which is not mandatory but good for references,
- a description of our case, and
- case tasks which is also not mandatory and can be added later on.

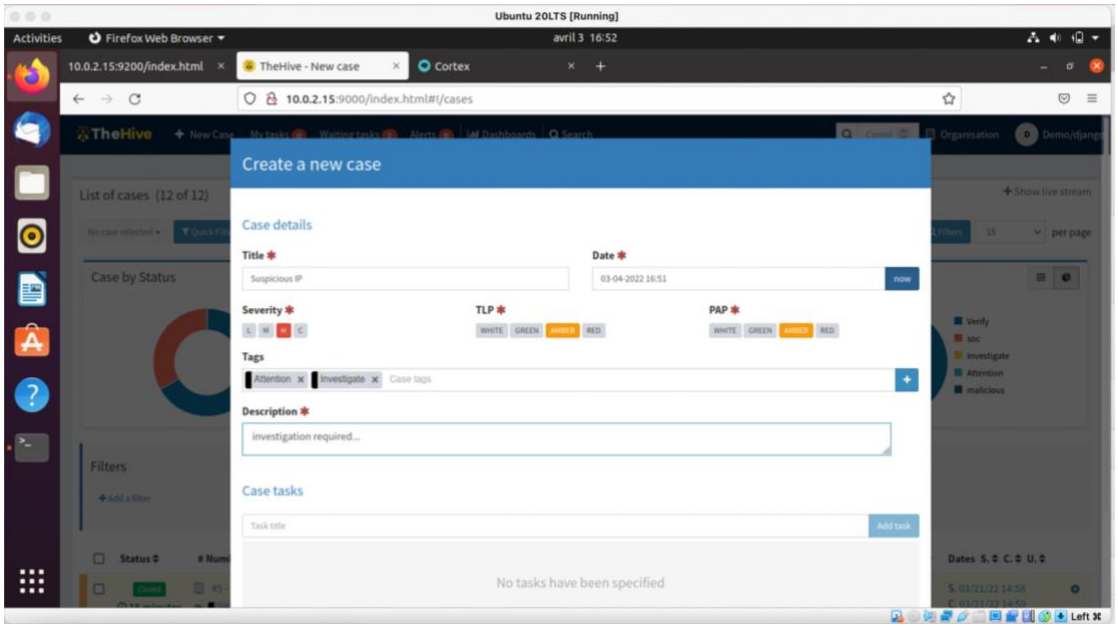

Figure 26 TheHive: providing new case data

We can assignee the case to another user by clicking on “Assignee” and selecting the user we want to assign it to.

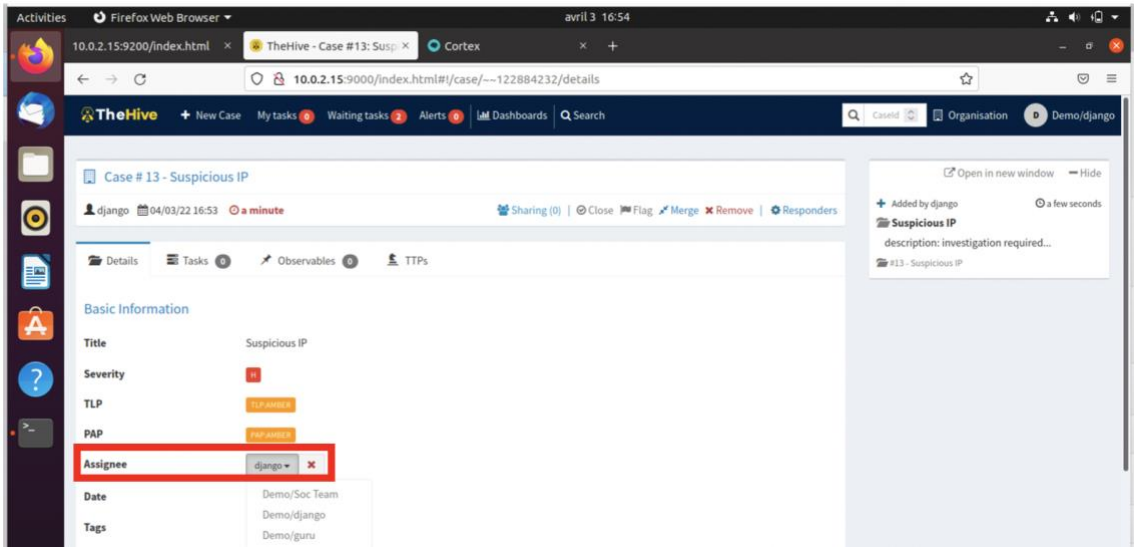

Figure 27 TheHive: viewing case information

- Creating tasks

We click on the “Tasks” tab to add a task to the case. We can provide a task group and a description. In this task, the task group is “soc” and the description of the task is “verify and blacklist”

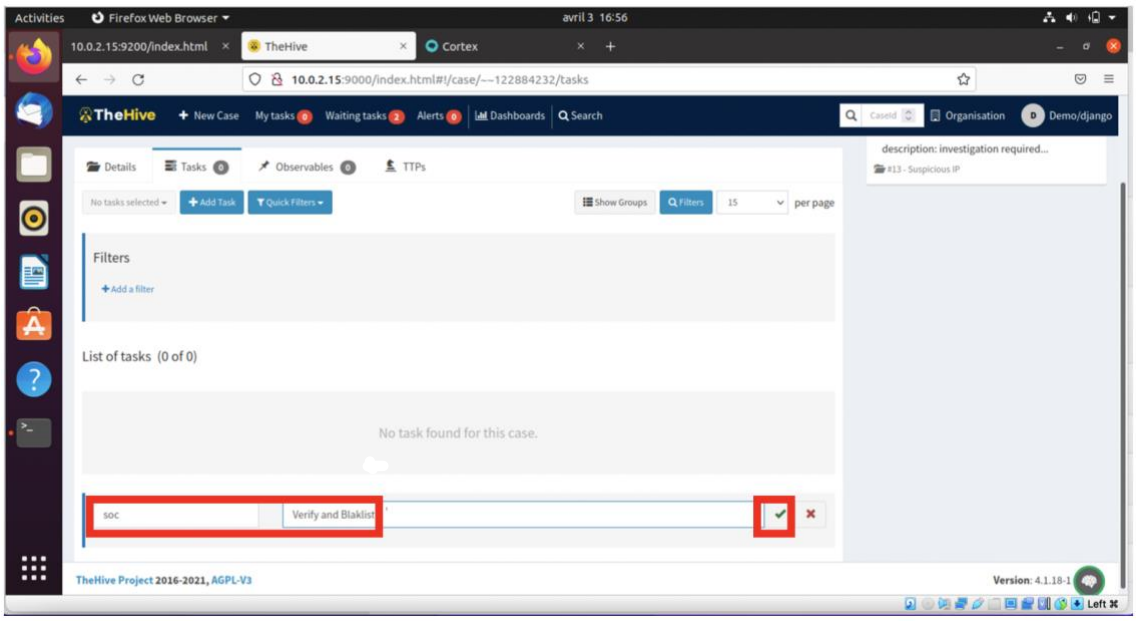

Figure 28 TheHive: creating a task

We add two tasks, namely “Check, verify and share info” and “Verify and blacklist” and we attributed them to different assignees.

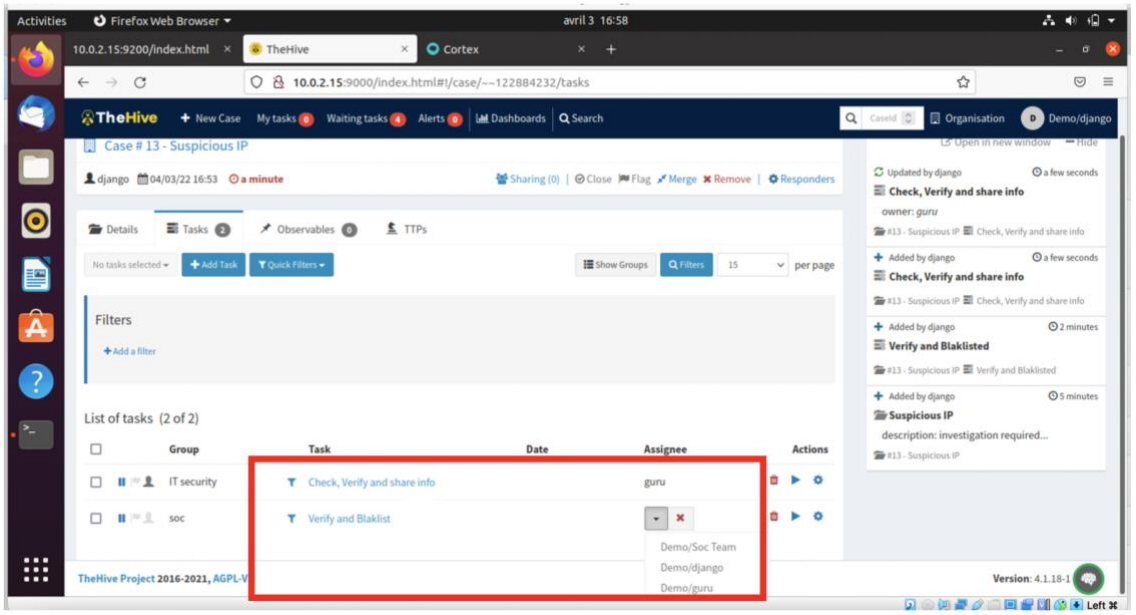

Figure 29 TheHive: assigning a task

- Running analysis

We then click on the “Observables” tab to add observables. A “Create new observable(s)” window opens. We have to provide:

- the observable types,
- the values, which is our observable,
- the TLP,
- tags, and
- a description of the observable.

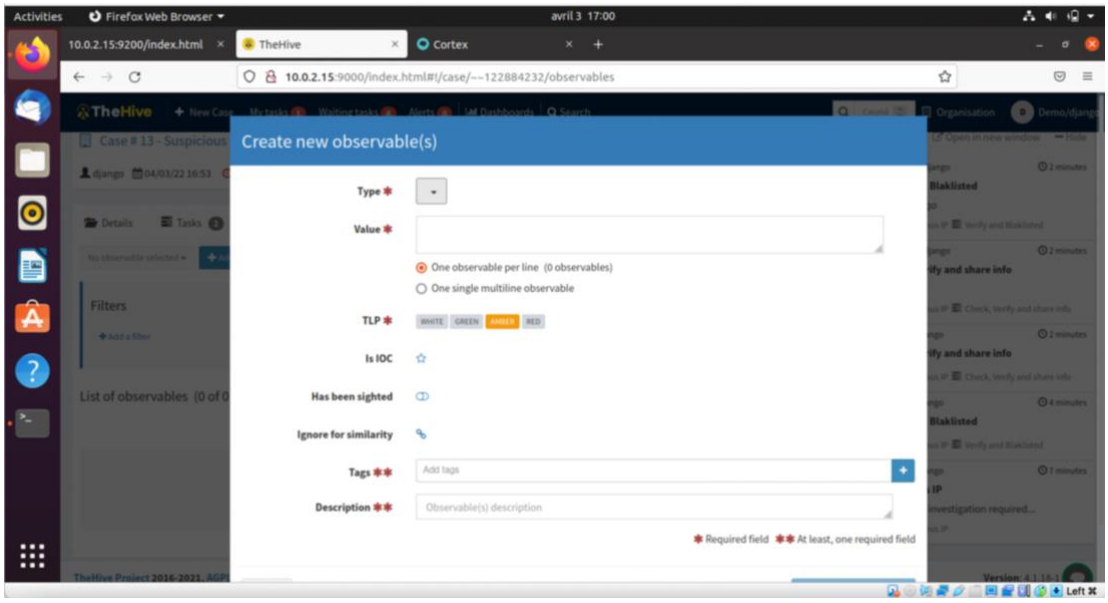

Figure 30 TheHive: create a new observable

In TheHive when an observable has been created, there is a possibility to see if the same observable has been used in other cases. To see if the observable has been used in other case, we can notice presence of an eye icon in the flags section. This eye icon shows that we have used the observable in at least one of our past cases. To see the details of the cases related to this observable, we click on the eye icon. This helps users to see what has been done before regarding the observable.

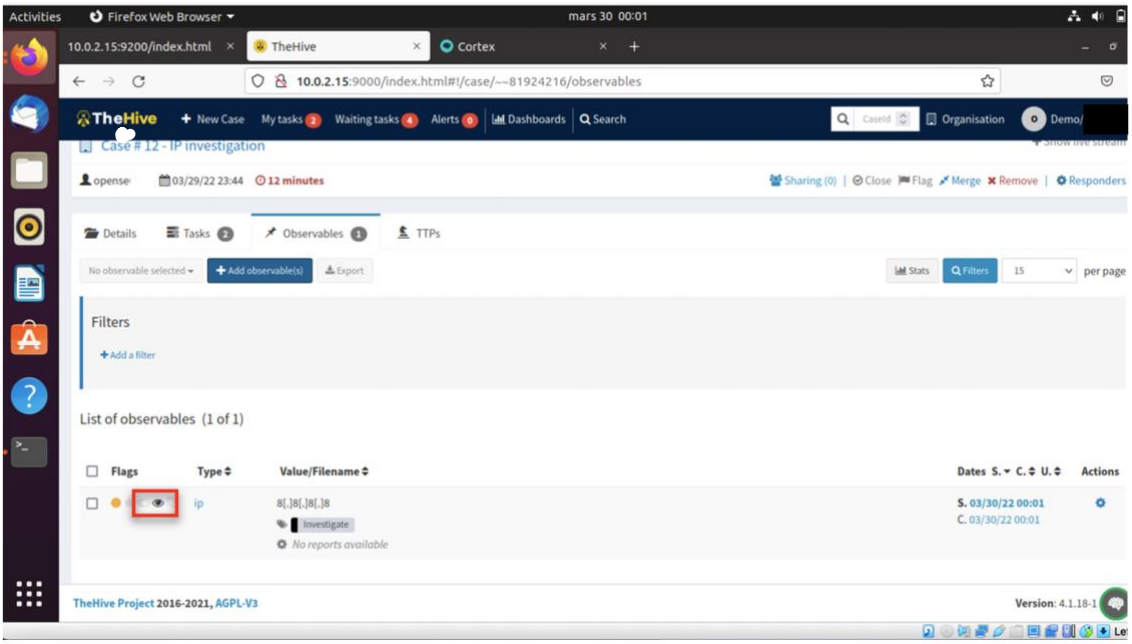

Figure 31 TheHive: viewing if used observable

In TheHive, we can have many observables for a single case. The “Guru” user from the “Demo” organisation can choose which observable he/she want to run analysers on. We choose one observable with the data type “IP”, and we click on “Run analysers”.

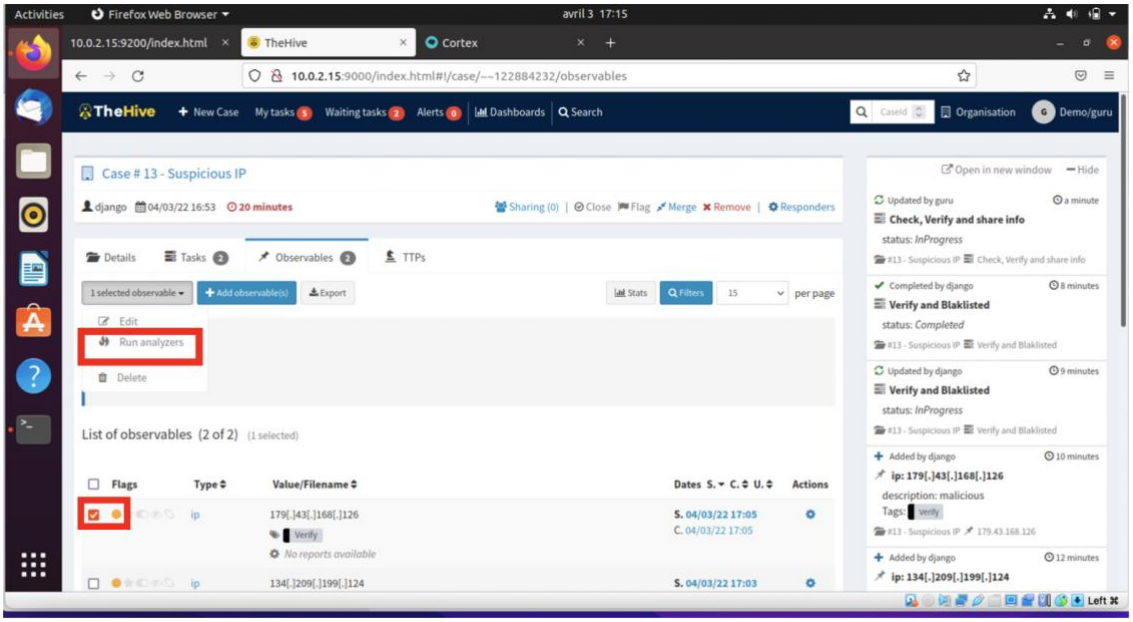

Figure 32 TheHive: running analysers

When the analysis complete, when can see its results. We can notice that the result has different colours.

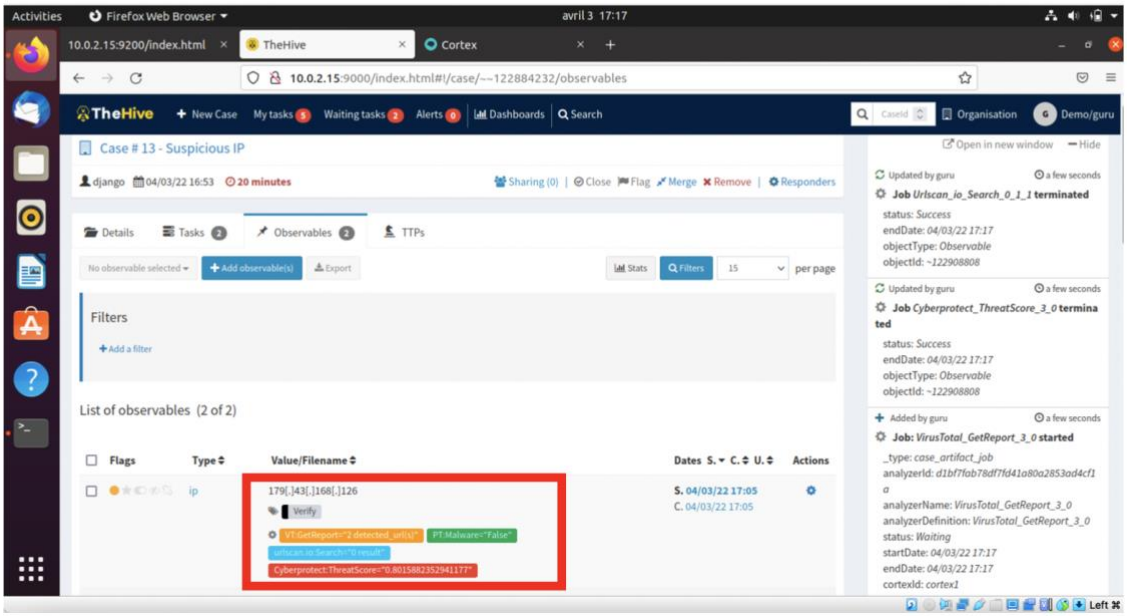

Figure 33 TheHive: viewing analysers' results

We click on the "CyberprotectThreatScore" in red to see the result of the analysis in details. We can see that this IP has a threat score of 80,2% and it has been identified as malicious.

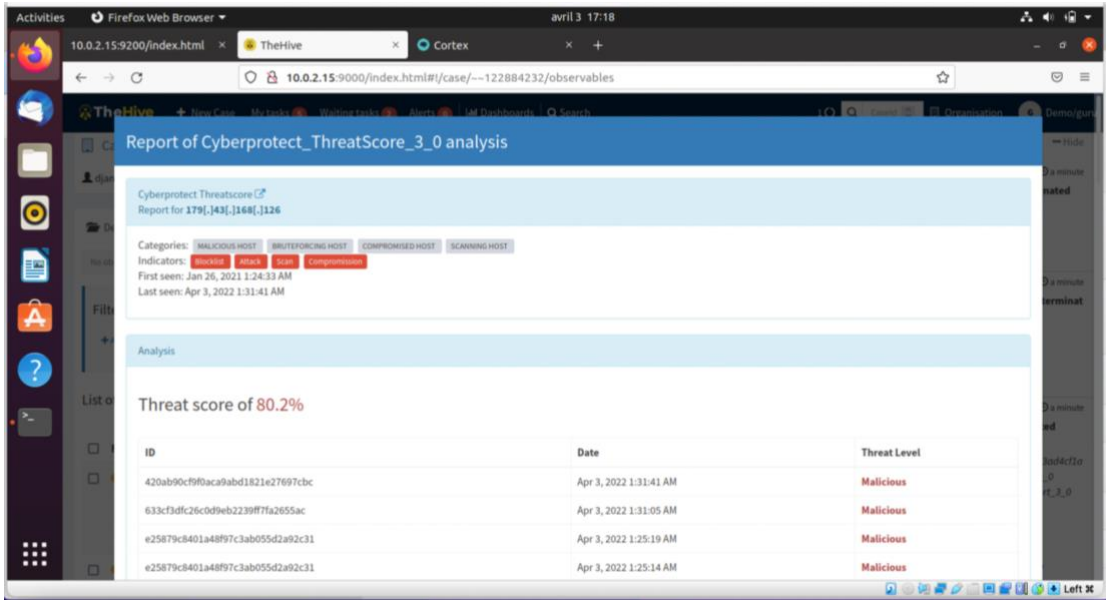

Figure 34 TheHive: viewing analysis reports

After this Guru can close the task. Guru selects the task he/she needs to close and click on "1 selected task" then on "Close".

Guru can also see the live feed for the Demo organisation. Indeed, every modification, update, and notification can be seen instantaneously for every user of an organisation. The user can also start a task by clicking TheHive live feed.

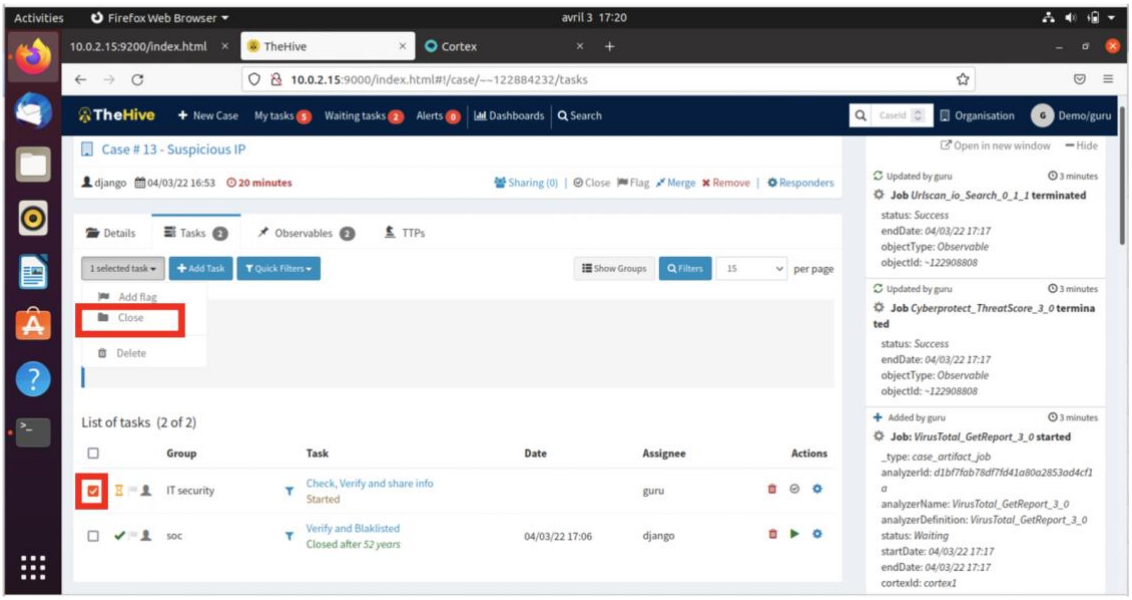

Figure 35 TheHive: closing a task

The user Guru clicks on the “Waiting tasks” tab on the upper ribbon and chooses on which task he/she wants to start an investigation. Guru can also see which tasks have been started, completed and are in progressed.

In the tasks section, Guru then runs an analyser against the observable provided. From the list of analysers Guru chooses four of them and runs them.

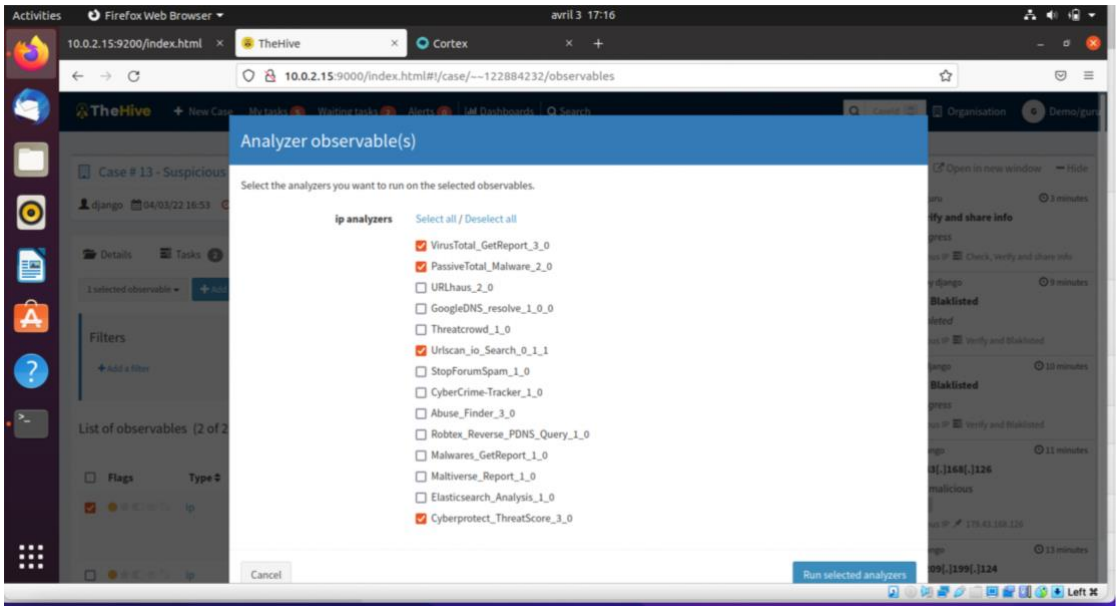

Figure 36 TheHive: choosing analysers

- Closing a case

When all tasks have been closed, we can close the case. To do so a user must have permission to manage a case. TheHive prompts us to indicate the status the case. We have to choose between:

- true positive, which means that the incident was initially identified as an attack and the investigation revealed that it was effectively an attack,
- a false positive, which means that a false alert was raised,
- an indeterminate, which means that there is not enough evidence to conclude positively or negatively on the status of the incident and this can be pushed to further investigation.

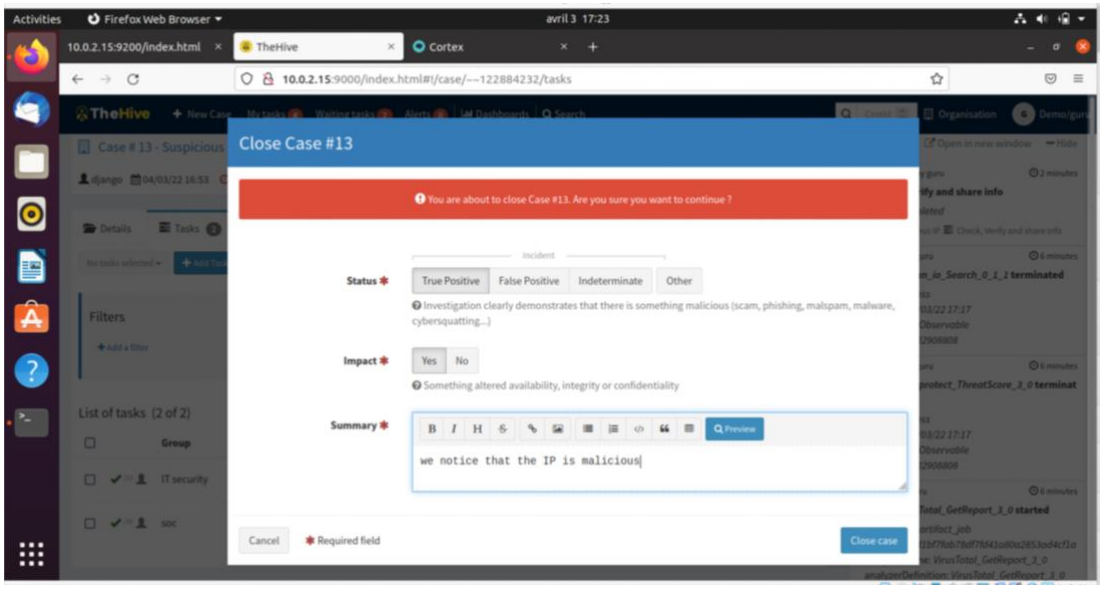

Figure 37 TheHive: closing cases TheHive

From the feed, the user Django can see when the tasks have been completed and closed by his colleague Guru.

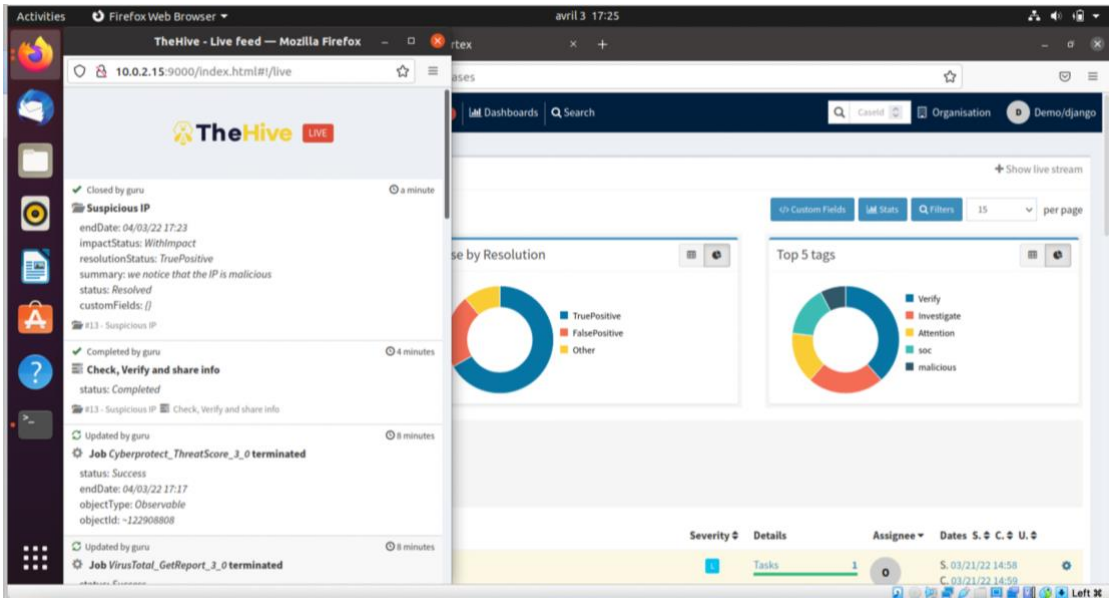

Figure 38 TheHive: checking status of closed case

- Running responder

After both tasks are completed and closed, Django can see in the results of the analysis that the IP address was malicious. He decides to take an action. This action is running a responder. Django clicks on “Run responders” and chooses “Wazuh_1_0”. This responder’s ability is to send information via email.

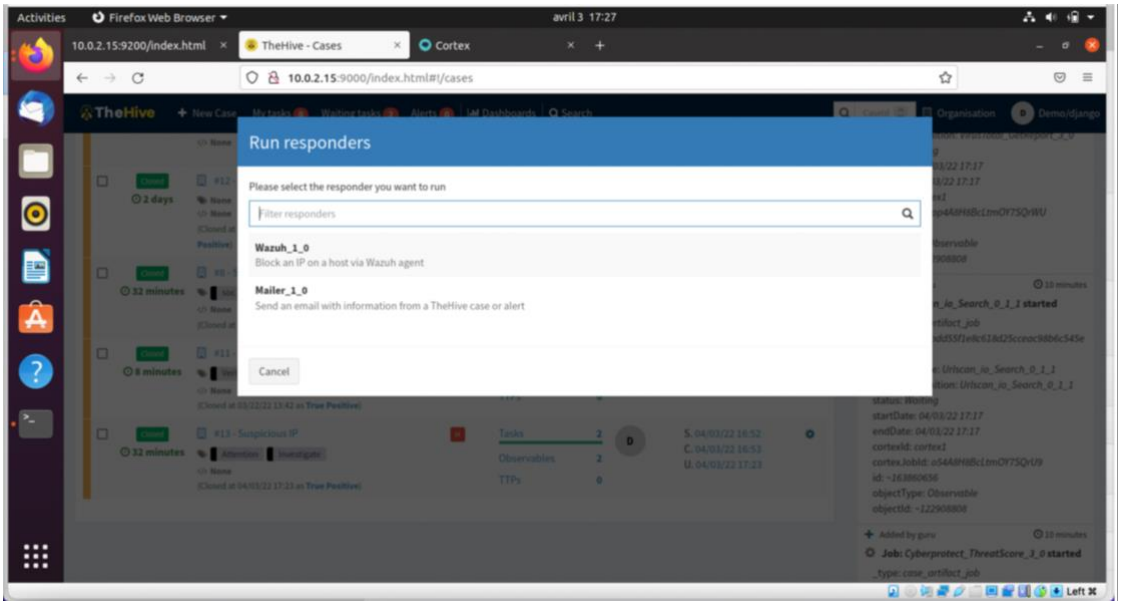

Figure 39 TheHive: running responders

- Requiring an action

TheHive gives the possibility of requiring an action on an open or closed task. This action can be assigned to a user by providing him/her with a log description that may include attachments such as pictures, documents, links, and others. It is also possible to reopen a task and share it.

Figure 40 TheHive: requiring an action (1/2)

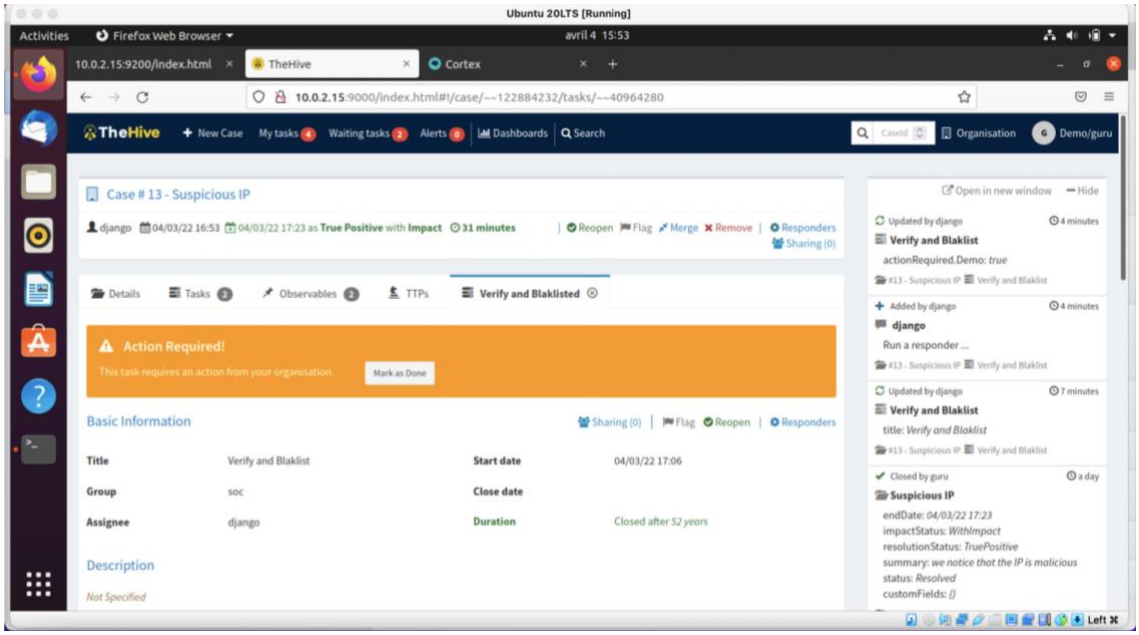

Figure 41 TheHive: requiring an action (2/2)

These two use-cases demonstrated that TheHive is effectively able to separate roles within an organisation thanks to its different roles available. The tool was also able to effectively manage the incident life cycle from the creation to the closure of a case.

6.3 Summary of findings

In this paper, we have seen that the power of Cortex resides in its ability to provide its users with access to a galaxy of analysers which are able to manage different types of observables. Indeed, if it is possible to go to the website of each analyser individually, it is time consuming for analysts to do so. By giving analysts the possibility to access all the analysers matching the type of the observable provided, Cortex greatly simplifies the work of the SOC team. The latter can keep track of all the analyses run inside a single application.

Integrating Cortex and MISP offers lots of insights. In Cortex, MISP is seen as an analyser which has the ability to detect whether an observable has already been spotted as IOC in MISP's large database of IOCs.

Combining Cortex and TheHive unlocks other functionalities of Cortex. It permits the analysis of multiple observables at once and also allows analysts to trigger active responses by making API calls to the responders' module of Cortex without leaving TheHive. All of this is made possible in an environment of collaboration between different analysts and organisations.

The capabilities of Cortex coupled with the other open-source tools TheHive and MISP could provide companies from critical sectors with the ability to identify threat information with the use of analysers and to share them within their organisations. Further work on the use of MISP to store and share the identified threat intelligence within the DYNAMO community is still to be done.

6.4 Comparison of results with objectives

In adequation with the set objectives, we were able to make list of Cortex's medium and high-level capabilities. We also provided a description of Cortex capabilities and its integration with MISP and TheHive. The description of the capabilities of Cortex when integrated with other open-source tools was essential as it made Cortex a more powerful tool. When compared with the CSIRT framework, we understood that Cortex was able to meet a restricted number of requirements from the framework due to its focus on observable analysis. We also were able to create a test environment for testing Cortex's capabilities and demonstrated these functionalities in this report. We suffered a setback when because of version incompatibility issues, we were not able to demonstrate the use of Cortex when integrated with MISP. We were, nevertheless, able to describe the functionalities linked to this integration in the previous section. Ultimately, this work can be used by the DYNAMO project to support the development of the module for the sharing of threat information of the task 4.2 in an environment of trust.

References

Electronic

Anand Groenewegen & Joris Janssen. 2021. TheHive Project: The maturity of an open-source Security Incident Response platform. Article from ResearchGate. Accessed 24 April 2022. https://www.researchgate.net/profile/Anand-Groenewegen/publication/352715439_TheHive_Project_The_maturity_of_an_open-source_Security_Incident_Response_platform/links/60ec1014b8c0d5588cef0c92/TheHive-Project-The-maturity-of-an-open-source-Security-Incident-Response-platform.pdf

CISA. Traffic Light Protocol (TLP) Definitions and Usage. https://www.cisa.gov/tlp#:~:text=The%20Traffic%20Light%20Protocol%20(TLP,by%20the%20recipient(s)

Davy Preuveneers & Wouter Joosen. 2012. Sharing Machine Learning Models as Indicators of Compromise for Cyber Threat Intelligence. Article from MPDI. Accessed 24 April 2022. https://www.mdpi.com/2624-800X/1/1/8/htm

Killcrece, G. Kossakowski, K-P. Ruefle, R. Stikvoort, D. West-Brown, M. & Zajicek, M. 2003. Handbook for Computer Security Incident Response Teams (CSIRTs). 2nd edition. Accessed 13 May 2022. https://apps.dtic.mil/sti/pdfs/ADA413778.pdf

Iklody, A. Dulaunoy, A. Wagener, G. & Wagner, C. 2016. MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform. Article from ACM Digital Library. Accessed 12 April 2022. https://dl-acm-org.nelli.laurea.fi/doi/pdf/10.1145/2994539.2994542

MISP Doc. 2021. Create an event based on a report. Accessed 28 March 2022. https://www.circl.lu/doc/misp/create-event-report/

MISP Doc. 2021. Taxonomies. Accessed 28 March 2022. https://www.circl.lu/doc/misp/taxonomy/

MISP GitHub. 2019. MISP taxonomies. Accessed 28 March 2022. https://github.com/MISP/misp-taxonomies/blob/main/PAP/machinetag.json

MISP Project. 2022. MISP training. Accessed 28 March 2022. https://www.misp-project.org/misp-training/misp-training.pdf

Nils Kuhnert. 2020. Palo Alto. Accessed 28 Match 2022. https://blog.thehive-project.org/tag/palo-alto/

ORACLE. 2022. Oracle VM VirtualBox. Accessed 28 Match 2022. https://www.oracle.com/virtualization/virtualbox/

Phishtank. No date. FAQ. Accessed 28 Match 2022. https://phishtank.org/faq.php

Paul Cichonski, Tom Millar, Tim Grance, Karen Scarfone. 2012. NIST Computer Security Incident Handling Guide. Accessed 13 May 2022. https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-61r2.pdf

Saâd Kadhi. 2017. TheHive, Cortex and MISP: How They All Fit Together. Accessed 13 May 2022. https://blog.thehive-project.org/2017/06/19/thehive-cortex-and-misp-how-they-all-fit-together/

Saâd Kadhi. 2019. TheHive 3.4.0 & Cortex 3.0.0 Released. Accessed 13 May 2022. https://blog.thehive-project.org/tag/support/

SonicWall. 2022. 2022 SonicWall Cyber Threat Report Key findings 02 Ransomware. Accessed 12 April 2022. https://www.sonicwall.com/2022-cyber-threat-report/?elqCampaignId=13998&sfc=7013h000000MiQZAA0&gclid=CjwKCAiAgbiQBhAHEiwAuQ6BkmbfNdHZWbIdJBPGBn4ut4T3yR5wDxM6JrGQbSMPEUk4O5ClyAmcVxoC7MsQAvD_BwE

TheHive-Project GitHub. 2020. Analysers requirement. Accessed 28 March 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/analyzer_requirements.md

TheHive-Project GitHub. 2020. How to create a responder. Accessed 28 March 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/api/how-to-create-a-responder.md

TheHive-Project GitHub. 2020. Install guide. Accessed 28 March 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/installation/install-guide.md

TheHive Project Documentation. 2021. Step-by-Step guide. Accessed 28 March 2022. https://docs.thehive-project.org/thehive/installation-and-configuration/installation/step-by-step-guide/

TheHive Project Documentation. 2022. Installation, operation and user guides. Accessed 25 March 2022. https://docs.thehive-project.org/thehive/

TheHive-Project. 2020. SECURITY INCIDENT RESPONSE FOR THE MASSES. Accessed 12 April 2022. https://thehive-project.org/

TheHive-Project GitHub. 2019. API Guide. Accessed 24 April 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/api/api-guide.md#job-model

TheHive-Project GitHub. 2019. How to Write and Submit an Analyzer. Accessed 24 April 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/api/how-to-create-an-analyzer.md

TheHive-Project GitHub. 2020. How to Write and Submit a Responder. Accessed 24 April 2022. https://github.com/TheHive-Project/CortexDocs/blob/master/api/how-to-create-a-responder.md

TheHive-Project GitHub. 2021. Cortex4py. Accessed 24 April 2022. https://github.com/TheHive-Project/Cortex4py

Virustotal. No date. How it works. Accessed 28 Match 2022.

https://support.virustotal.com/hc/en-us/articles/115002126889-How-it-works

Wazuh Docs. 2022. Elasticsearch single-node cluster. Accessed 28 March 2022. https://documentation.wazuh.com/current/installation-guide/more-installation-alternatives/elastic-stack/distributed-deployment/step-by-step-installation/elasticsearch-cluster/elasticsearch-single-node-cluster.html

Unpublished

DYNAMO. 2021. Dynamic Resilience Assessment Method including a combined Business Continuity Management and Cyber Threat Intelligence solution for Critical Sectors. Unpublished.

Figures

Appendices

Appendix 1: List of medium and high-level capabilities of Cortex

Related to users

- Creating a user,
- Updating a user,
- Listing users,
- Seeing details about a user and,
- Deleting a user.

Related to organisations

- Creating an organisation,
- Updating an organisation,
- Listing organisations,
- Seeing details about an organisation and,
- Deleting an organisation.

Related to analysers and responders

- Enabling/Disabling and configuring analysers and responders to be used by Cortex,
- Analysing observables using analysers,
- Responding to an alert or an investigation using responders via application such as TheHive.

Related to jobs from analysers

- Listing all jobs from analysers,
- Seeing reports from jobs from analysers and,
- Deleting jobs.

Appendix 2: Assessment of Cortex according to the CSIRT framework

Type of requirement	Requirement	Is Cortex able to satisfy (Y/N)	How is the requirement satisfied?	Grade
Vulnerability and artifact analysis	(1) Using results artifacts analysis	Y	Cortex performs artifact analysis and helps analysts in their analysis of the incidents	5/5
Vulnerability and artifact analysis	(2) Using results of vulnerability analysis	N		
incident life cycle handling	(3) Notification of impacted stakeholders when closing an incident ideally as part of the conversation ongoing during the incident	N	Need to assess capabilities of TheHive, SIRP instead	
incident life cycle handling	(4) Reusage of the same tracking number when an incident is reopened and a new one if separate matter	N	Need to assess capabilities of TheHive, SIRP instead	
incident life cycle handling	(5) Possibility to mark incidents as related and possibility to merge them	N	Need to assess capabilities of TheHive, SIRP instead	
incident life cycle handling	(6) Ability to keep track of information related to the incident	N	Need to assess capabilities of TheHive, SIRP instead	
incident analysis	(7) the ability to categorise incidents using multiple categories before the team receives it	N		
incident analysis	(8) receiving the logs in a manner that is appropriate to its categories	N		
incident analysis	(9) having the ability to authenticate the log	N		
incident analysis	(10) cleaning the logs from sensitive information not necessary for the analysis	N		
incident analysis	(11) the ability to send extract of the log as incident follow-up to the relevant audience	N	Need to assess capabilities of TheHive, SIRP instead	
incident analysis	(12) perform some profiling for networks and systems and recognize a normal behaviour	N		
incident analysis	(13) Create a policy that determines how long logs should be kept	N		
incident analysis	(14) Correlate events between one another and synchronize host clocks	N		
incident analysis	(15) Keep a “knowledge base of information”	Y	Thanks to the use of analysers, Cortex provide access to the database of scanning tools containing a lot of information able to help analysts in the incident analysis.	4/5
incident analysis	(16) Using search engines on the internet to research information about potentially abnormal activity	Y	The use of Cortex’s analysers can be a substitute to the use of search engines in the case where there is a potential indicator of compromise	3/5
incident analysis	(17) Use packet sniffers to gather more data	N	Other tools such as Wireshark could potentially be used	

incident analysis	(18) Have the ability to filter data by category of indicators	N		
incident analysis	(19) Request help from others	Y	The use of analysers provide access to analyses done by other communities such as the MISP community and can in certain cases replace the need to request help from experts.	3/5
triage	(20) Assigning and using tracking numbers	N	Need to assess capabilities of TheHive, SIRP instead	
triage	(21) Registering the contact information	N	Need to assess capabilities of TheHive, SIRP instead	
announcement	(22) elements that trigger announcement should be clearly defined	N	Need to assess capabilities of TheHive, SIRP instead	
announcement	(23) categorisation that helps understanding which kind of announcement will be needed	N	Need to assess capabilities of TheHive, SIRP instead	
announcement	(24) Prioritisation of announcements to ensure that the message has the right impact on its audience.	N	Need to assess capabilities of TheHive, SIRP instead	
announcement	(25) published information should be cleared and	N	Need to assess capabilities of TheHive, SIRP instead	
announcement	(26) Consideration of the channels for distributions depending on the audience and the sensitivity of the information	N	Need to assess capabilities of TheHive, SIRP instead	