



A coordinated framework for cyber resilient supply chain systems over complex ICT infrastructures

## D5.1 IT-1 FISHY release integrated

Document Identification			
Status	Final	Due Date	30/11/2021
Version	1.0	Submission Date	21/01/2022

Related WP	WP5	Document Reference	D5.1
Related Deliverable(s)	D2.2	Dissemination Level (*)	PU
Lead Participant	TID	Lead Author	Jose Manuel Manjón
Contributors	UPC, UMINHO, UC3M, POLITO, STS, TUBS AND XLAB	Reviewers	ATOS
			TUBS

Keywords:
Sandbox, FISHY platform, Data, Security, Virtualization.

This document is issued within the frame and for the purpose of the FISHY project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 952644. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission.

This document and its content are the property of the FISHY Consortium. All rights relevant to this document are determined by the applicable laws. Access to this document does not grant any right or license on the document or its contents. This document or its contents are not to be used or treated in any manner inconsistent with the rights or interests of the FISHY Consortium or the Partners detriment and are not to be disclosed externally without prior written consent from the FISHY Partners.

Each FISHY Partner may use this document in conformity with the FISHY Consortium Grant Agreement provisions.

(\*) Dissemination level: **PU**: Public, fully open, e.g. web; **CO**: Confidential, restricted under conditions set out in Model Grant Agreement; **CI**: Classified, **Int** = Internal Working Document, information as referred to in Commission Decision 2001/844/EC.

## Document Information

List of Contributors	
Name	Partner
Jose M. Manjón Cáliz	TID
Diego R. López	TID
Eva Marín Tordera	UPC
Henrique Santos	UMINHO
Luis Félix González Blázquez	UC3M
Cataldo Basile	POLITO
Daniele Canavese	POLITO
Grigorios Kalogiannis	STS
Mounir Bensalem	TUBS
Jan Antič	XLAB

Document History			
Version	Date	Change editors	Changes
0.1	03/12/2021	TID	Index
0.2	15/12/2021	TID	Contribution to sections 1 and 2
0.3	17/12/2021	UPC, UMinho	Contribution to section 3
0.4	21/12/2021	UC3M	Contribution to sections 2 and 6
0.5	22/12/2021	POLITO	Contribution to section 4
0.6	17/01/2022	TUBS	Contribution to section 5
0.7	19/01/2022	XLAB	Contribution to section 3
0.8	19/01/2022	TID	Version ready for internal review
0.9	19/01/2022	ATOS, TUBS	Internal review comments
1.0	21/01/2022	ATOS	FINAL VERSION TO BE SUBMITTED

Quality Control		
Role	Who (Partner short name)	Approval Date
Deliverable leader	Jose Manuel Manjón (TID)	21/01/2022
Quality manager	Juan Alonso (ATOS)	21/01/2022
Project Coordinator	Antonio Alvarez (ATOS)	21/01/2022

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated	<b>Page:</b>	2 of 20
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU
	<b>Version:</b>	1.0	<b>Status:</b>
			Final

# Table of Contents

Document Information .....	2
Table of Contents .....	3
List of Tables.....	4
List of Figures .....	5
List of Acronyms .....	6
1 Introduction .....	7
1.1 Purpose of the document .....	7
1.2 Relation to another project work .....	7
1.3 Structure of the document .....	7
2 Sandbox platform.....	8
2.1 Architecture.....	8
2.2 Installation.....	9
2.3 Hello World example .....	10
3 Integrated Trust Manager .....	11
3.1 T3.1- Security & Privacy Data Space Infrastructure (SPI) .....	11
3.2 T3.2- Trust & Incident Manager (TIM) .....	11
3.3 Trust Manager .....	11
4 Integrated Security and Certification Manager .....	13
5 Intent-based Resilience Orchestrator and Dashboard .....	15
6 Secure Infrastructure Abstraction .....	17
7 Conclusions .....	19
References .....	20

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated			<b>Page:</b>	3 of 20
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0
				<b>Status:</b>	Final

---

## List of Tables

---

*Table 1. Plan for Integration*.....12

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	4 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

---

## List of Figures

---

*Figure 1. Sandbox architecture.....8*

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	5 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## List of Acronyms

Abbreviation / acronym	Description
AC&IdM	Access Control and Identity Management
D5.1	Deliverable number 1 belonging to WP5
DM	Data Manager
CLI	Command Line Interface
CNL	Controlled Natural Language
EC	European Commission
EDC	Enforcement and Dynamic Configuration
F2F	Farm-to-Fork
HPL	High-level Policy Language
IRO	Intent-based Resilience Orchestrator
K8s	Kubernetes
NED	Network Edge Device
NLP	Natural Language Processing
PMEM	Predictive Maintenance Monitoring
SACM	Security Assurance & Certification Management
SeCM	Security Capability Model
SIA	Secure Infrastructure Abstraction
SPI	Security & Privacy Data Space Infrastructure
TIM	Trust & Incident Manager
VNF	Virtual Network Functions
VM	Virtual Machine
WP	Work Package

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated	<b>Page:</b>	6 of 20				
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

---

# 1 Introduction

---

## 1.1 Purpose of the document

---

The objective of this deliverable is to report the progress of the FISHY platform implementation and present the first version of it. Here is described the FISHY platform, integrated on a Sandbox. The idea is that this Sandbox will contain all the FISHY components and provide communication between them.

Regarding other tasks, this document makes reference to the activities that have been carried out in other work packages that will provide their tools into the FISHY platform, with the definition of the different interfaces, components and all the necessary inputs to make the FISHY platform work.

## 1.2 Relation to another project work

---

All the integration outputs of the work packages involved on the design and development of the different tools are reflected in this document. The WP3 integrates the Trust Manager module, the WP4 integrates the Security and Certification Manager and regarding this WP5, the tasks T5.1 and T5.2 contribute with the Intent-based Resilient Orchestrator & Dashboard and the Secure Infrastructure Abstraction, respectively.

## 1.3 Structure of the document

---

This document is structured in 5 major chapters:

- **Chapter 2** presents the FISHY platform, the Sandbox. Here is explained the architecture of the platform, the process of installation and how to execute the *Hello World* example.
- **Chapter 3** presents the work done on the integration of the Trust Manager (WP3).
- **Chapter 4** presents the work done on the integration of the Security and Certification Manager (WP4).
- **Chapter 5** presents the work done on the Intent-based Resilience Orchestrator and the Dashboard (T5.1).
- **Chapter 6** presents the work done on the Secure Infrastructure Abstraction (T5.2).

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	7 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## 2 Sandbox platform

This section describes what is the FISHY Sandbox, its architectural design, the installation process, and the description of the hello-world example to showcase how to use the Sandbox.

All the information regarding the FISHY Sandbox that will be described in this section of the document is also available in the link below. In this case, the information is showcased using a video, showing the architecture of the first iteration of the FISHY Sandbox, the installation process, and the deployment of the hello-world example [12].

### 2.1 Architecture

Figure 1 presents the design of the FISHY Sandbox. The purpose of this Sandbox is to provide a virtual environment capable of supporting the execution of FISHY components and other relevant functions, such as Virtualized Network Functions (VNFs) developed during the project lifetime.

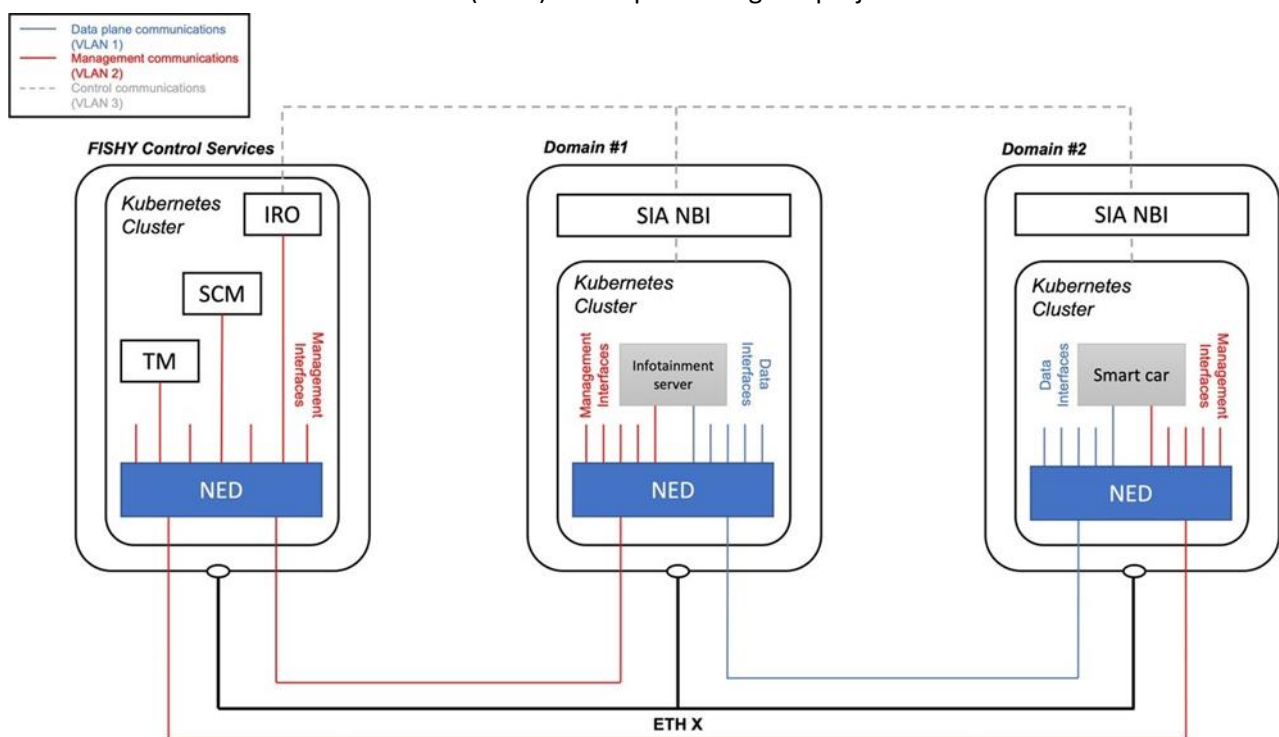


Figure 1. Sandbox architecture.

As it can be seen in the picture, the first version of the Sandbox is provided as a set of different domains. One of the domains hosts the FISHY control services (e.g., the Trust & Incident manager, the Security and Certification Management, and the Intent-based Resilience Orchestrator & Dashboard). Two additional domains provide the abstraction of a Network Function Virtualization Infrastructure (NFVI). This way, the Sandbox design is flexible to support different testing requirements, involving a single organization domain, or even several organization domains if needed.

To support inter-domain communications, the Sandbox includes a functional Network Edge Device (NED) at every domain. NEDs establish all the management communications between fishy control

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated	<b>Page:</b>	8 of 20
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU
<b>Version:</b>	1.0	<b>Status:</b>	Final



services and other FISHY components, such as VNFs (management communications are shown in red color in the figure). They also enable inter-domain data-plane communications between VNFs (data-plane communications are shown in blue color). In the first version of the sandbox (IT-1), the NED offers a set of pre-created management and data-plane interfaces, providing secure network connectivity to applications and modules in the FISHY Sandbox.

The first software release of the Sandbox is based on Kubernetes (K8s) [1]. Each domain is provided as a K8s cluster that supports the deployment of all the modules and applications in the Sandbox and includes a representation of an ICT infrastructure so that any pilot can be deployed on a domain. In turn, each K8s cluster is packaged into a virtual machine. This way, the Sandbox can be downloaded by interested partners as a set of three interconnected virtual machines. Still, its design is flexible to accommodate different virtual infrastructure management solutions, e.g., OpenStack [2]. A NED implementation can be also made available for OpenStack deployments. In addition, an up-to-date version of the Sandbox is hosted in the 5TONIC laboratories in Madrid, to support test activities and serve as a reference implementation for use case development.

## 2.2 Installation

This section of the deliverable covers the installation of the IT-1 version of the FISHY Sandbox. A full description of the installation process can also be found in the repository of the project. In this repository, all the images required to perform the installation are available to be downloaded.

After completing this installation process, a FISHY Sandbox infrastructure like the one present in Figure 1 will be available for its use in the environment where it has been deployed. Each machine in the Sandbox has a NED element deployed in each domain, with the same network interfaces and IP tunnels shown in the figure.

The basic prerequisites to deploy and operate with the Sandbox platform are the following:

- Three Virtual Machines (VMs) using the “*fishy-sandbox-baseline.qcow*” image. Each VM must have at least 2CPUs and 2GBs of RAM for its execution.
- All machines must be interconnected through a virtual network that provides Internet connectivity. Each VM must have one network interface connected to this virtual network
- The deployment of each VM with a single network interface is recommended. If more interfaces are aggregated to the VM, the script will attach the NED to the interface used to reach the internet (i.e., the default route). Therefore, the VM must be able to reach the rest using this interface.

The installation process that must be followed to set up the Sandbox is the following:

1. Login into the machine using the following credentials:
  - *User*: admin-fishy
  - *Password*: admin-fishy
2. If not present in the VM, download the configuration file(publicly accessible):

```
wget https://github.com/Networks-it-uc3m/FISHY-Sandbox-development/blob/main/sandbox-config/sandbox-config.bash && chmod +x sandbox-config.bash
```
3. Start the installation process using the following command. If prompted, introduce the password of the VM:

```
./sandbox-config.bash
```
4. If the VM to be configured is the fishy-control-services host, introduce in the command line the “y” character:

```
Is this machine fishy-control-services?[y/n]
```

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	9 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

y

Otherwise, type “n” and write in the command line the corresponding domain that the VM will represent (“domain-1” or “domain-2”):

```
Is this machine fishy-control-services?[y/n
```

```
n
```

```
Is this domain-1 or domain-2?[domain-1/domain-2]
```

```
domain-1
```

5. Introduce the respective IP addresses of the other VMs when prompted by the console:

```
Please, enter Domain 1 IP:
```

```
10.0.0.1
```

```
Please, enter Domain 2 IP:
```

```
10.0.0.2
```

6. Wait for the process to be completed. If asked by the command line through the following message, write the “y” character and press Enter:

```
cp: overwrite '/home/ubuntu/.kube/config'?
```

```
y
```

7. Once the following message appears, the VM is ready to be used:

```
Node [fishy-control-services/domain-1/domain-2] ready!
```

You can also check that it is properly installed by introducing the command `kubectl get pods` and checking that the “`ned-[DOMAIN-NAME]`” is up and in the “`Running`” status.

## 2.3 Hello World example

---

This example mimics the deployment of a simple use case in the FISHY platform, where an infotainment server in Domain-1 will communicate with a smart car located in Domain-2. All components have a management interface that the fishy-control-services domain will use to communicate with them (for management purposes).

The detailed example can be found on the GitHub page (publicly accessible)[7].

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	10 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

---

## 3 Integrated Trust Manager

---

### 3.1 T3.1- Security & Privacy Data Space Infrastructure (SPI)

---

The SPI comprises two main modules: the Data Manager (DM), where data transformation (including sanitization) and dispatching (based on RabbitMQ) takes place; and the Access Control and Identity Management (AC&IdM), aiming to enforce authentication and authorization rules. In IT-1, the DM reads log files from the infrastructure, transforms them into CEF format, and sends them to the appropriate RabbitMQ channel, adopting the classification defined by the metrics taxonomy (still under development). The other FISHY modules read data from the RabbitMQ by subscribing to the required channel(s). The template of the code developed for testing purposes is available at [8].

The AC&IdM module exposes a service (URI) through which i) applications previously registered (TIM and SCM) are authenticated and require an access token to have access to any other FISHY compliant module, and ii) front end components request user authentication (not required in IT-1). The application code used for testing purposes, including the token access request and RabbitMQ access data, is available at [9].

### 3.2 T3.2- Trust & Incident Manager (TIM)

---

The TIM component is composed of several modules, Vulnerability Assessment, Incident Detection, Impact Assessment and Mitigation. Sharing data between these modules is facilitated by the Threat/Attack Repository, which is used for storage of incoming metrics and the results of analysis performed by the various modules of TIM. The Threat/Attack Repository also serves as the point of integration with other platform modules such as the IRO and EDC, that will perform further actions based on the results of TIM. The communication between modules, both internal to TIM and other platform modules, is facilitated by an HTTP(S) REST API, that allows all CRUD (Create, Read, Update, Delete) operations over the data entities handled by TIM; and a publish/subscribe RabbitMQ message bus, that allows any involved components to receive real-time updates on any relevant new or updated data.

The Swagger (OpenAPI v2) definition of the Threat/Attack Repository is available at [10] and an example of a subscriber component, a websocket server providing data updates to a webUI, is available at [11].

### 3.3 Trust Manager

---

For IT-1 implementation, Trust Manager (TM) uses a simplified workflow for communication and interaction between the lower infrastructure (SIA), SPI and TIM; the SPI is currently not yet FISHY agents to get direct data from the use cases infrastructure. These tools will use the Rabbit MQ protocol for getting the data from the infrastructure. In the current implementation, only Wazuh working in TIM is right now using the Rabbit MQ protocol to do the communication with the F2F use-case to get the log files. The identity Manager module in SPI is using the keycloak tool to handle the access to the Central Repository in TIM. In the proposed workflow for IT-1, all the tools working in TIM will get the access token from keycloak to access the Central repository. Currently, we did not test this functionality because the work is in progress. After obtaining the token, the different tools working in the TIM will use REST API to write the data in the central repository.

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	11 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

**Table 1. Plan for Integration**

<b>Sr N°</b>	<b>Task</b>	<b>Responsible Body</b>	<b>Status</b>	<b>Deadline</b>
1	Integration T3.1 - SPI in Sandbox	UPC/UMihno	In Progress	15/02/2022
2	Integration T3.2 - TIM in Sandbox	UPC/XLAB	In Progress	15/02/2022
3	Finishing PMEM	UPC	In Progress	15/02/2022
4	Design and implementation of the central repository	UPC/XLAB	In Progress	15/02/2022
5	Integration of Complete Trust Manager (TM)	UPC/XLAB/UMinho	In Progress	28/02/2022

## 4 Integrated Security and Certification Manager

The two main components of the WP4 are the EDC (in charge of refining and enforcing the policies) and the SACM (in charge of producing the security certifications). Moreover, WP4 is in charge for defining several data models used at cross-WP level, i.e., the representations policies (at different abstraction levels) and the security capability.

The Security Capability Model (SeCM) and the workflow to manage it have been developed and are in a stable version. SeCM provides support for the description of the security features of a variety of security controls (e.g., generic packet filter, generic VPN gateways, iptables, XFRM). The SeCM is also used to automatically provide also the Medium-Level policy representation, an abstract language for configuring the capabilities available at one specific NSF which is automatically derived from the NSF capabilities.

An initial prototype of the Register and Planner component of the EDC implementing the capability model has been developed as an extension module (http-service) of the BaseX XML database. This service reads the information about all the NSFs available in one FISHY domain which is provided in form of an XML Catalogue. The Register and Planner provides an interface to perform queries on the catalogue, to identify NSFs that own specific features, to compare the security capabilities owned by two NSFs, etc. The integration in the sandbox of the Register and Planner is scheduled for the early months of 2022.

The Translation features of the Enforcer component of the EDC have been also defined in the framework of the SeCM. This component is currently a stand-alone Java-based tool. It will be released as a service and integrate in the sandbox during the Q1/Q2 or 2022.

In addition, the Enforcer will interface with the SIA (by sending the configurations to the various NSFs and requesting the current landscape) and WP4 will start its integration once the SIA APIs are defined and stable.

In addition, WP4 partners are actively collaborating with WP5 ones for finalizing the high-level policy language (HPL) schema in order to be used as target for the translation of the intents. The initial model of the HPL has been heavily extended to describe new fields and to support a new category of requirements, the reactive policies needed by the use cases, and for implementing the features in WP3 and WP5.

Finally, WP4 has also defined an initial version of a remediation framework that will be used in conjunction with the mitigation block defined in WP3. This component provides a remediation recipe language and an interpreter that allows the Controller module of the EDC to propose and implement mitigations. The threat intelligence modules, implemented in WP3, will leverage the mitigation capabilities of the TIM that, in turn, will make use of this recipe language to instruct the EDC on how to react to an anomalous situation and attacks.

The Security Assurance & Certification Management, under Task 4.2 is responsible for providing a real time, continuous assessment of the security posture of the complex ICT systems inside FISHY project and it is enabled by a purpose-built Evidence Collection Engine, which aggregates the required evidence from multiple sources related to the operation of individual components, as well as the overarching processes where these components are involved in. This functional group of modules also includes an Audit and Certification subcomponents, leveraging the evidence-based approach of the Assurance solution integrated into the platform.

Currently both major components of SACM have been successful tested by integrating them with SYNELIXIS premises during Q3 of Y1, using our primary versions. The latter have been exploited and

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	13 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

---

developed furthermore and currently are provided to the FiSHY platform as dockerized – containerized systems.

Integration of the SACM with the sandbox has been started in late Y1 and will be continued and finalized during the Q1 of Y2.

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	14 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## 5 Intent-based Resilience Orchestrator and Dashboard

The task T5.1 is composed by two main components, which are the Intent-based Resilience Orchestrator (IRO) and the Dashboard. IRO aims at automating the interactions between the user defining high level intents and the system applying high level policies, using natural language processing (NLP) and AI techniques. The Dashboard component is an interface that integrates all the dashboards used among other FISHY components. The workflow and interactions between IRO and other components and inside IRO have been defined and a first prototype of IRO has been developed. Two flows are defined: A top-down flow where a user (network administrator) interacts with IRO via an IRO-Dashboard in order to provide commands in the form of high-level intents, after that IRO pre-process and translates the intents to configure the required policies, and finally the policies are verified and forwarded to the network controller represented by the EDC components. In the bottom-up flow, the data gathered by monitoring tools defined in TIM component is used by IRO in order to react on network events and reconfigure the status of used policies using EDC component. Moreover, IRO considers the human in the loop and interacts with the user via the Dashboard in order to confirm and verify the concluded policies that are aimed to be applied.

IRO is composed by several blocks, which are the IRO-Dashboard/IRO-CLI, the intent manager, the policy configurator, the intent compiler, Learning & reasoning and the knowledge base. IRO-Dashboard or IRO-CLI is a frontend interface used by the user to input the intents to IRO, where few CLI commands have been developed in the current version (functionalities: add intent, get status, reset status, push intents).

Using NLP techniques, the intent manager translates the input text in the intent into a structured format containing several fields such as the intent type and parameters. At the current version a controlled natural language (CNL) is used to define a specific grammar for few intents, regular expressions are currently used to extract the information from the text, and we aim at extending and enriching the language to consider more language expressions, and to use smarter techniques of information retrieval. The policy configurator is responsible for matching the extracted requirements from the input intent with exploitable policies, where it configures the required policies stored in the knowledge base to satisfy the user requirements. At the current version, the policy configurator has a predefined method for each type of intents, where the possible configuration scenarios are analyzed in advance, and policy templates that can be reconfigured are defined stored. IRO can be extended and enriched with functionalities simply by adding policy templates (where the policy has to be adapted to the capability model used in EDC component), and a matching method. Such feature will be analyzed and defined in the next iteration.

The intent compiler component receives the configured policies and verify their validities with the initial intents, considering that some conflicts between intents might occur, which needs to be solved before enforcing the policies. The learning & reasoning is a component that receives information from monitoring tools (TIM), checks if an alert or recommendation has to be sent to the user via the Dashboard, and then reacts based on that. The learning and reasoning can automatically react to alerts when a predefined intent has been given by the user, where it triggers the policy configurator to configure policies and pass them to EDC using the intent compiler. This component has been defined but not yet been developed. Finally, the knowledge base component is used to save the intent structure, which contains a dictionary of vocabularies that can be used in the intent translation, as well as the existing policy templates. In the current prototype, we are using Elasticsearch in order to save the intent structure in a JSON format.

The integration of the current IRO version on the Sandbox is already achieved, where IRO software contains two containers: IRO source code and Elasticsearch. They can leverage the NED to

<b>Document name:</b>	D5.1 IT-1 FISHY release integrated			<b>Page:</b>	15 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

communicate with the container hosted on the other VM of the Sandbox. In the next step, we aim at defining the interface with the other components such as EDC and TIM. For the use cases, we plan to analyze each use case separately and extract few possible intents that the user may provide, prepare hello world example and develop it. Currently an example from F2F use case has been analyzed in collaboration with WP4 leader, which required an enhancement in the capability model. The next step is to develop the examples and plan how to extend it to other use cases.

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	16 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final



## 6 Secure Infrastructure Abstraction

The Secure Infrastructure Abstraction (SIA) module offers a northbound interface to the other blocks and components of the FiSHY platform. The SIA northbound interface provides an abstract and technology agnostic view of the NFV infrastructure resources available at an organization domain. It supports the management and orchestration of network services and VNFs making use of that infrastructure resources. A full description of the SIA, its architecture and functionality is already available in D2.2[3]. However, for the sake of completeness, this section will include a summary of the most relevant aspects of the SIA, as well as the development status of its components and the future for its integration in the FiSHY platform.

In the design and implementation of the SIA northbound Application Programming Interface (API), the consortium will consider the following functions within the scope of an organization domain: a) management of NFV descriptors (e.g., upload/delete/update) that define network services, Virtual Network Functions (VNFs), and Physical Network Functions (PNFs); b) lifecycle management of network services that operate on the organization domain; c) collection of performance information regarding the execution of network service instances; d) issuing notifications under fault conditions (e.g., a failure on a virtual network that impacts the connectivity of a network service or VNF); and e) provision of information on the capacity of NFV infrastructure resources within the organization domain.

On the other hand, to keep compatibility with relevant standardization efforts on NFV, the preliminary work on the SIA design considers the possibility that NFV descriptors follow the YANG models defined by ETSI [4]. In addition, the design work envisions a SIA northbound API aligned with the RESTful API specification defined by ETSI to support the interaction between an Operation/Business Support System (OSS/BSS) and an NFV orchestrator [4].

The development of the SIA northbound interface is currently being considered for implementation in IT-2. In the current version, IT-1, a set of functionalities that will be performed by the northbound interface are currently present in the Sandbox using the Kubernetes interface available in each one of the FiSHY domains. Nevertheless, this approach is limited, as it does not support the automated establishment of link-layer connectivity between virtual functions (Kubernetes pods), as well as the management of Kubernetes services making use of this connectivity. To aid automation, a first step in the realization of the SIA northbound API contemplates the development of a Kubernetes wrapper that supports the management (creation/update/deletion) of virtual networks on Kubernetes clusters.

It is important to observe that the SIA module must provide the API functions regardless of the underlying NFV management and orchestration technologies (e.g., Open-Source MANO [5], OpenStack or Kubernetes) and SDN controllers (e.g., OpenDaylight, Ryu, etc.) that may be in use within a FiSHY domain. To this purpose, the design of this module considers the utilization of an adaptable southbound interface. The SIA southbound interface will support the interaction with specific management and orchestration software stacks and SDN controllers that exist in a domain, through the utilization of pluggable translating modules (plug-ins). This plug-in-based approach decouples the design and the implementation of the SIA module from specific NFV management and orchestration technologies, which can be supported as add-ons to the southbound interface.

To perform its functionalities, the SIA uses a set of Network Edge Devices (NEDs) deployed in every FiSHY domain. In principle, this component oversees the provision of the switching functionality necessary to enforce forwarding rules for the traffic received onto the point-to-point links built between NED peers located in different domains. Further information can be found in D2.2, where the NED and its functionalities are described in further detail.

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated			<b>Page:</b>	17 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

IT-1 provides the first iteration of the NED component, which builds IP tunnels seen in Figure 2 between each NED through Virtual eXtensible Local Area Networks (VXLANS) [6]. These NEDs allow applications, deployed as Kubernetes pods, to communicate between each other using point-to-point links. Each application can rely on the own Kubernetes interface to perform the rest of functionalities that do not involve communications through the NEDs (e.g., Internet access). Their deployment is automatically performed once the Sandbox domains are installed using the available script shown in previous sections of this document.

For the NED that will be introduced in IT-2, the automatic configuration of the applications attached to the NED will be considered. In other words, applications will be able to automatically attach to the NEDs in their respective domains and configure their networking interfaces without any manual configurations needed to be performed (in contrast on how it is done so far, for example requiring to manually assign IP addresses in the containers of an application). Furthermore, the next iteration of the NED will support more characteristics of Kubernetes deployments, including replicas for service scalability, Kubernetes service compatibilities, etc.

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	18 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

---

## 7 Conclusions

---

The main conclusion of this deliverable is that the first version of the Sandbox has been released and the partners can download it to test their tools locally to be further integrated. Also, a description of the work packages and tasks involved on the integration have been made in order to make a follow-up of this work.

Next steps regarding integration includes the improvement of the Sandbox by automatization the deployment of the services and automatic IP assignment and also continue with the integration of all the components and tool of each use case.

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated				<b>Page:</b>	19 of 20	
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final

## References

- [1] The Linux Foundation. “Kubernetes: Production-Grade container orchestration”. Accessed December 17, 2021. [Online] <https://kubernetes.io>
- [2] Open Infrastructure Foundation. “OpenStack: The most widely deployed open source software in the world”. Accessed November 25, 2021. [Online] <https://www.openstack.org>
- [3] FiSHY. D2.2 – IT-1 architectural requirements and design. 2021.
- [4] ETSI Group Specification NFV-SOL 006, Network Functions Virtualisation (NFV) Release 3; Protocols and Data Models; NFV descriptors based on YANG Specification, version 3.3.1, August 2020.
- [5] ETSI OSM. “Open Source Mano”. Accessed December 17, 2021. [Online] <https://osm.etsi.org>
- [6] M. Mahalingam, D. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridhar, M. Bursell, C. Wright, Virtual eXtensible LocalArea Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks, RFC 7348, 2014, doi:10.17487/RFC7348.
- [7] <https://github.com/Networks-it-uc3m/FiSHY-Sandbox-development/tree/main/Guides#hello-world>
- [8] <https://github.com/hdsantos/FiSHY-SPI/blob/main/FiSHY-prod-ex.py>
- [9] <https://github.com/hdsantos/FiSHY-SPI/blob/main/FiSHY-cons-ex.py>
- [10] <https://github.com/H2020-FiSHY/TIM/blob/develop/tar/api/swagger/swagger.yaml>
- [11] <https://github.com/H2020-FiSHY/TIM/blob/develop/ws-notifier/server.js>
- [12] <https://vm-images.netcom.it.uc3m.es/FiSHY/videos/FiSHY-Sandbox-installation.mp4>

<b>Document name:</b>	D5.1 IT-1 FiSHY release integrated			<b>Page:</b>	20 of 20		
<b>Reference:</b>	D5.1	<b>Dissemination:</b>	PU	<b>Version:</b>	1.0	<b>Status:</b>	Final