



Foundations of Trustworthy AI – Integrating Reasoning, Learning and Optimization
TAILOR
Grant Agreement Number 952215

SRIR Development in Collaboration with the ICT-48 CSA, PPP-AI, and AI-on-demand Platform
Report

Document type (nature)	Report
Deliverable No	11.5
Work package number(s)	11
Date	M18
Responsible Beneficiary	UNIBO, ID #10
Author(s)	Michela Milano
Publicity level	Public
Short description	Coordination of TAILOR Strategic Research and Innovation Roadmap activities with the AI-on-Demand Platform, PPP, and CSA initiatives

History			
Revision	Date	Modification	Author
Version 1	4 April 2023	-	Michela Milano

Document Review		
Reviewer	Partner ID / Acronym	Date of report approval
Marc Schoenauer	3 - INRIA	11 April 2023
Fredrik Heintz	1 - LiU	11 April 2023

This document is a public report. However, the information herein is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Table of Contents.

Summary of the report	2
Organisation	2
1. Purpose of the SRIR	3
2. Gaps on the current version of the SRIR	4
3. Second internal consultation	5
4. External consultation	5
4.1 EU initiatives connected to TAILOR	6
4.2 Process of engagement	9
4.3 Questions and expected contributions	10

Summary of the report

This report describes the process and the strategy for involving strategic European AI initiatives in the definition of the TAILOR Strategic Research and Innovation Roadmap, version 2.

For creating a cohesive ecosystem, it is important that all initiatives collaborate and coordinate to create roadmaps that cover different aspects, but also are aligned.

In the report, we recall the objectives of the TAILOR SRIR, identify the current gaps and shape the process to improve the SRIR by acting internally and externally to the TAILOR project with the most related initiatives and outline a process for their involvement.

Organisation

The following people have been involved in the Deliverable:

Partner ID / Acronym	Name	Role
10 - UNIBO	Michela Milano	Responsible SRIR
3 - INRIA	Marc Schoenauer	Responsible WP2

1. Purpose of the SRIR

TAILOR is a Network of Excellence devoted to creating an ecosystem and delivering solid research results around Trustworthy AI, focusing on Learning, Optimization and Reasoning. In work package 2, TAILOR will develop a Strategic Research and Innovation Roadmap (SRIR). Its purpose is not only to shape the main research directions of the project, but it goes beyond the project duration with the purpose of influencing and providing guidelines to the whole European AI community as well as to the European Commission in future work programmes. The ambition of the roadmap is to outline the foundations of Trustworthy AI for the years 2020-2030.

The three objectives of the SRIR, aligned with the ones of TAILOR, are the following:

- OB1: Strengthen and enlarge a pan-European network of excellent researchers on the Foundations of Trustworthy AI
- OB2: Strengthen the scientific foundations for Trustworthy AI, and translate them into technical requirements to be adopted broadly by industry.
- OB3: Strengthen collaborations between academic, industrial, governmental, and community stakeholders on the Foundations of Trustworthy AI.

The TAILOR project is clearly the core source of content for the SRIR. The scientific work-packages

- WP3 Trustworthy AI
- WP4 Integrating AI Paradigms and Representations
- WP5 Deciding and Learning How to Act
- WP6 Learning and Reasoning in Social Contexts
- WP7 Automated AI

contribute to define the current and long-term research directions by analysing the current gaps and needs, mainly from a research community perspective, but also from an industrial viewpoint. By acting internally, we can improve the current version of the SRIR by (1) defining actions and priorities and (2) connecting more tightly the Learning, Optimization and Reasoning (LOR) aspect of the project with the Trustworthy AI (TAI) dimensions. At the moment, the two components of the project are somehow connected but we should investigate and plan more deeply what LOR can do for enforcing TAI dimensions like for example fairness, privacy, explainability, robustness, sustainability and, on the other hand, how to make LOR techniques trustworthy.

In addition, to spread the concept of Trustworthy AI pervasively in Europe, both in academia and industry, it is essential to connect the SRIR development with the other initiatives that

form the European AI ecosystem. We identify the main initiatives that we consider for the liaison with TAILOR on the SRIR topic and then we define the process to engage them in the SRIR definition and continuous update.

The deliverable is organised as follows. We first analyse the gaps that are present in the TAILOR SRIR, then we trace the path for closing these gaps by consulting the TAILOR consortium, with particular attention to scientific task leaders, but trying to involve all partners in the collection of input and contributions. After this internal consultation, we open our perspective to other EU initiatives and for each of them we identify preliminary questions that can be addressed by external communities. Then we define the process of engagement with specific actions, meetings and tools that we plan to use to collect feedback on the current material and input of new material for the SRIR.

2. Gaps on the current version of the SRIR

The TAILOR project has delivered the first version of the SRIR and produced a public document to be shared inside and outside the boundaries of the project. The public version is available on the TAILOR web site¹. The first version has been written by consolidating the input coming from the TAILOR consortium, and in particular by collecting content from the EREB, composed by scientific WP leaders.

Still, many gaps need to be filled in for the second version:

Merging TAI and LOR: The dimensions of trustworthy AI and the Learning, Optimization and Reasoning side of the TAILOR project are, in the SRIR, somehow disconnected. It is important that they are more glued and that we identify directions for achieving Trustworthy AI features of AI systems using LOR and, on the other way round, to use LOR for enforcing TAI features. This is crucial for the overall success of TAILOR and could represent an important funding direction for the EC in the following work programs. For this gap, it is important to involve the TAILOR consortium to identify these connections and define research topics aimed at strengthening the connection between TAI and LOR.

Priorities and actions: While the first version of the SRIR identified a list of research directions that need to be investigated to produce innovation and to bring ahead the foundational research in all scientific WPs, it is still open how these directions should be prioritised, and which actions can be put in place for achieving them. For this gap, it is important to involve the TAILOR consortium to identify short- and long-term objectives and actions.

Connections with other initiatives: The first version of the SRIR has been produced by consulting the TAILOR consortium, while Trustworthy AI and Learning Optimization and Reasoning could in principle affect and be affected by other European Strategic initiatives related to AI and related to other disciplines. For this reason, we need to first identify initiatives that we want to involve in the SRIR definition process and to shape a timeline for their involvement.

¹ <https://tailor-network.eu/wp-content/uploads/2022/07/TAILOR-Roadmap-Full-Version-1.pdf>

3. Second internal consultation

For improving the connection between Trustworthy AI and Learning, Optimization and Reasoning, we plan to have two steps: one online, and one in-person during the TAILOR conference in Siena. For the first one, we will use MURAL, which is an online collaboration tool that is particularly suited to brainstorming, by organising content into rooms and folders for an easier consolidation.

The main questions we want to ask to the TAILOR consortium are the following:

- Which are the main gaps to create AI systems merging learning, optimization and reasoning that are also explainable, fair, robust, accountable, secure, privacy preserving and sustainable?
- How can learning, optimization and reasoning help achieve trustworthiness in AI systems?
- How can the research directions that are highlighted in the first version of the SRIR be prioritised? Which are the most urgent? For which directions we are closer to a solution and for which direction we need more research and experiment?
- Is there any direction that has emerged in the last year that deserves attention? For instance, the current discussion about the block of ChatGPT in some countries, or the letter to stop experiments on LLM for six months.
- Do connections with the AI Act need attention from the TAILOR perspective? Which are the articles of the AI Act connected to TAI dimensions and how do they require expertise on the technical side?
- Is the attention of the industrial world toward Trustworthy AI aspects high enough? Is there any action that the TAILOR project could do or propose to increase the awareness and interest toward TAI aspects?

We plan to have a meeting in May 2023 with the TAILOR consortium and then consolidate the inputs received before the TAILOR conference. During the TAILOR conference, we plan to have a presentation of the SRIR and ask participants to provide feedback.

4. External consultation

The second version of the SRIR aims at collecting inputs from European Strategic initiatives that are relevant to TAILOR. First, we identify the initiatives connected, and we define questions to be posed to each of them and also concrete actions to involve them.

4.1 EU initiatives connected to TAILOR

The European Commission has contributed to shape an Artificial Intelligence ecosystem where several initiatives coexist and cooperate. Figure 1, borrowed from the WP4 of the AI4EU project, shows the current main AI initiatives. In the following, we describe the ones we consider more related with TAILOR, therefore they are the initiatives we need to interact with to obtain input for the SRIR and influence their further updates to better take into account, if necessary, the trustworthiness of AI, that might not be their primary concern.

ICT48 Networks of Excellence and the associated CSA VISION: These are the closest initiatives with TAILOR, as TAILOR is indeed one of the ICT48 Networks of Excellence. The others are ELISE, AI4Media and Humane-AI-Net. They are Research and Innovation actions devoted to developing networks across Europe and therefore create the conditions for European foundational research, with low TLR. The coordination and support action VISION coordinates all four ICT48 Networks. Each of them has to produce a Strategic Research and Innovation Roadmap, one of the Work Packages of VISION is to synchronise, align, and synthesise the different roadmaps under the same conceptual and operational strategic canvas.

AI-on-demand platform and follow up projects: The AI-on-demand platform is a strategic initiative funded by the European Commission through the first project AI4EU, that has provided the first investment to create the platform. AI4EU finished in December 2021. The platform is live on <https://www.ai4europe.eu/> and contains several tools, an Ethical observatory, and an experimentation playground.

Six ICT49 projects have then been funded under Horizon 2020 to enrich the platform either through additional services or verticals. These ICT49 are StairwAI², AIPlan4EU³, AI4Copernicus⁴, DIH4AI⁵, Bonsapps⁶ and I-nergy⁷. These projects are funded under the innovation action instrument. These projects do not have to produce a roadmap, but it is important to interact with them to obtain input for the TAILOR roadmap.

Finally, also in Horizon Europe ICT49, a CSA called AI4Europe⁸ has been funded with the aim of developing the platform for the research community, while another project will be funded under Digital Europe to develop the platform strategy toward the industrial ecosystem and drive it to the market. Until July 2023 a preparation action for this Digital Europe project has been funded under the name Pre-Pai⁹. It will be important as well to liaise with these initiatives.

² <https://stairwai.nws.cs.unibo.it/>

³ <https://www.aiplan4eu-project.eu/>

⁴ <https://ai4copernicus-project.eu/>

⁵ <https://www.dih4ai.eu/>

⁶ <https://bonsapps.eu/>

⁷ <https://i-nergy.eu/>

⁸ <https://www.ai4europe.eu/>

⁹ Being a preparatory action the project does not have a web site.

AI, data and Robotics PPP and related CSA The AI, Data and Robotics Partnership¹⁰ is one of the European Partnerships in digital, industry and space in Horizon Europe. To deliver the greatest benefit to Europe from AI, Data and Robotics, this Partnership will drive innovation, acceptance, and uptake of these three technologies.

The Partnership will boost new markets, applications and attract investment, to create technical, economic, and societal value for business, citizens, and the environment. By 2030, European sovereignty is expected in the development and deployment of trustworthy, safe and robust AI, Data and Robotics, compatible with EU values and regulations. The PPP has developed a Strategic Research, Innovation and Deployment Agenda and it would be important to connect with the PPP to obtain important input about the TAILOR SRIR. Connected associations and initiatives are of course relevant to the TAILOR SRIR, starting with the private partner of the PPP, Adra (AI, Data and Robotics Association), and its founding members BDVA (now DAIRO), euRobotics, CLAIRE, ELLIS, and EurAI.

Furthermore, a CSA called Adra-e¹¹ has been funded to support the AI, Data and Robotic partnership, create a lively ecosystem, and support the Adra association.

HPC and Cybersecurity related initiatives Despite being less related to the core topics of TAILOR, initiatives related to High Performance Computing, and Cybersecurity are clearly relevant for TAILOR. In particular, for the HPC side, we will connect either with the EuroHPC¹² Joint Undertaking or with PRACE¹³ depending on their availability to collaborate. For the link cybersecurity we plan to connect with ECSO¹⁴.

As a disclaimer, we have to say that while for AI related initiatives, the collaboration spirit is high and well supported by the EC so we are very confident that we obtain feedback and answer to questions, with other research and innovation communities the risk of not receiving feedback is higher. For this reason, we will identify two or three key persons of these initiatives and interview them either together or one by one. For HPC we will identify three key persons working at CINECA and BSC that are hosting two pre-exascale machines funded by the EC and one member of PRACE. For cybersecurity we will try to reach some key person of the European Cyber Security Organization.

¹⁰ <https://adr-association.eu/>

¹¹ <https://adra-e.eu/>

¹² <https://eurohpc-ju.europa.eu>

¹³ <https://prace-ri.eu/>

¹⁴ <https://ecs-org.eu/>

The European AI Landscape

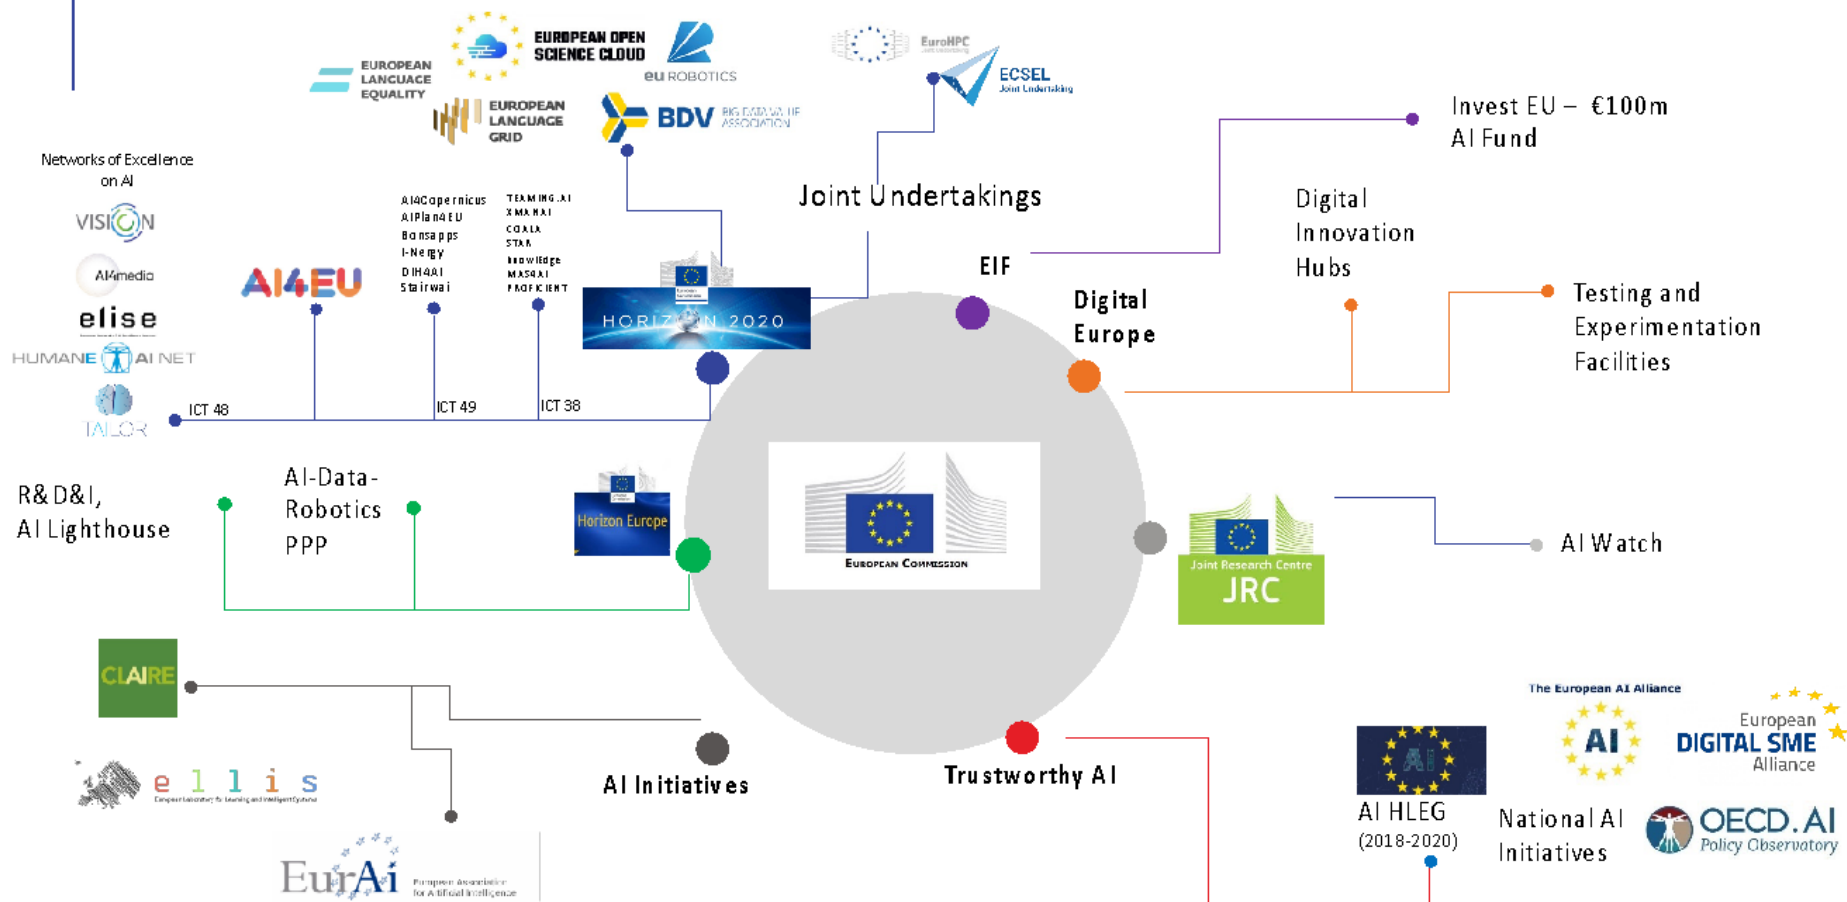


Figure 1: the European AI ecosystem

4.2 Process of engagement

Overview: Once we have defined the initiatives that are targeted toward the engagement with TAILOR, we have to define the WHO, WHAT and WHEN. WHO is the TAILOR responsible that will take care of the action, WHAT concerns the action we have to deliver and WHEN is the timing.

Concerning the timing, we have in TAILOR WP2 decided that the first version of the SRIR will be merely based on inputs from the TAILOR consortium, while for the second version we plan to engage with all the above-mentioned initiatives. Starting from the end of March 2023, when the first version of the SRIR is delivered, we plan to interact with the other initiatives following the process defined in Figure 2.

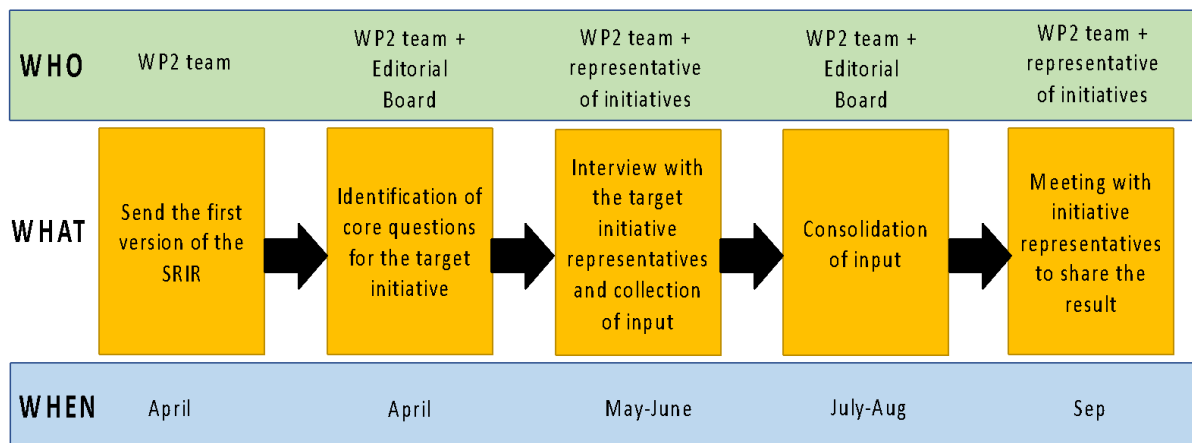


Figure 2: Process for engaging related initiatives

Starting from April 2023, the SRIR Editorial Board will send the first version of the SRIR to the initiatives identified in Section 4.1. For each target initiative, the WP2 team and the SRIR editorial board should identify questions that are specific for the initiative itself. For each initiative, in May-June we will run interviews/workshops with selected representatives to collect input for the SRIR. An important event that will help us collect results from AI-related initiatives is the AI Ecosystem Forum promoted by the AI4Copernicus Project that will be held in Athens on June 29-30. All the main AI initiatives will be present, and we plan to run some on-line meetings before, as done for the internal consultation so that we can have a draft agenda for the AI ecosystem forum.

Then the WP2 team and the SRIR Editorial Board will consolidate the inputs obtained by the interviews and then we share the final result, and we start disseminating the final version of the TAILOR SRIR document.

4.3 Questions and expected contributions

We have identified some initial questions that could be posed to each initiative thus focussing on the peculiarities of the initiative and its links with TAILOR. The Editorial Board will be involved to enrich the set of questions or change them if needed.

ICT48 Networks of Excellence and the associated CSA VISION

Networks of Excellence are mainly interested in developing fundamental research. They cover aspects like human-machine interaction, machine learning and deep learning, AI for media with particular attention to fake everything and generative AI. It is clear that each of them has to consider aspects related to trust. For this reason, the main questions we would pose are the following:

- Which are the main open challenges related to trustworthiness in your specific domain? Are they somehow already investigated from a practical perspective? Are there solutions already or are they still open research fields?
- Will your specific domain be influenced by AI regulations?
- Is there any connection with learning, optimization and reasoning and their integration thereof that would improve your domain and its trustworthiness aspects?
- Which are the most urgent aspects that you would consider, and which are more longer terms. Which actions can be put in place to address these trustworthiness aspects?

Also in this case, as for the internal consultation, we plan to reach out to various projects (in a meeting all together) and then share the results during the TAILOR conference that will also include VISION representatives.

AI-on-demand platform and follow up projects

Concerning the AI-on-demand platform, the AI4Europe project has to shape a Strategic Roadmap of which the Trustworthy AI part is a chapter. There is an intersection between the editorial board of TAILOR and the one of AI4Europe. Therefore, we will join forces to collect feedback on common aspects as all projects are devoted to reach the same objectives of making the community aware and pushing the research in the direction of TAI dimensions. For what concerns instead ICT-49 projects, they are very much organised around TAI related matters as there is a Trustworthy AI working group up and running, with regular meetings.

The main questions we would pose to this audience will be the following:

- Which is the role that the AI-on-demand platform could play for promoting Trustworthy AI to both the research community and the innovation ecosystem?

- Can the AI-on-demand platform be considered as a unique entry point for experimenting TAI dimensions in AI systems?
- Could the AI-on-demand platform provide an experimentation environment for measuring the trustworthiness of an AI solution?
- Could the AI-on-demand platform provide a sort of “certification” of trustworthiness that could label AI systems? Should this certification also identify issues related to data?

For these issues we will first contact the Trustworthy AI working group of ICT49 and then share the results during the Athens meeting at the end of June. For what concerns the AI4Europe project we will run a joint event to collect feedback on TAI.

AI, data and Robotics PPP and related CSA

The AI, Data and Robotics PPP is devoted to addressing the research and innovation ecosystem at the intersection of AI, data and robotics. Therefore, the questions we will ask them are of course related to their field of expertise.

- What are the main fields involving at least two of the three fields (AI, Data, and Robotics) that need attention from a trustworthiness perspective? Are they still open challenges or somehow solved problems?
- What are the main impacts on trustworthiness that come from data? Do they need a separate treatment for obtaining a sort of certification or can they be handled together with the rest of the AI life cycle? How does this relate to the AI Trust Label that Adra-e will define?
- Which are the most urgent aspects that you would consider, and which are more longer terms. Which actions can be put in place to address these trustworthiness aspects?
- Do you think that companies engaged by the PPP pay enough attention to TAI aspects? And what about researchers?

For this initiative, we need to organise two separate meetings one for the input collection and possibly one for the consolidated input return. We plan to do these meetings one in May and one in September.

HPC and Cybersecurity related initiatives

Concerning HPC, IoT and Cybersecurity we can discuss issues that are more high level and not specifically connected with AI development. These are some questions that we could propose:

- TAI dimensions require the development and the use of AI systems to be more computationally expensive: for example providing explanations for an optimization problem output would require additional inference at search time; data augmentation

for ensuring fairness would require additional storage. Which are the main areas you see for which HPC could help the development and use of TAI systems?

- Security is one of the dimensions of TAI systems. On one hand the system should be safe and on the other hand it should be robust and avoid causing unintentional harm. Do you foresee any research challenge that is still open and that needs to be addressed urgently?

For these initiatives, we need to organise two separate meetings one for the input collection and possibly one for the consolidated input return. We plan to do these meetings one in June and one in September possibly in form of interviews.