SAFEXPLAIN

Safe and Explainable
Critical Embedded Systems based on AI

D6.6 Initial Exploitation Report

Version 1.0

Documentation Information

Contract Number	101069595
Project Website	www.safexplain.eu
Contratual Deadline	30.09.2023
Dissemination Level	PU
Nature	R
Authors	Carlo Donzella (EXI), Luigi Dematteis (EXI)
Contributors	all partners
Reviewer	Janine Gehrig Lux (BSC)
Keywords	exploitation, IPs, IPRs, copyright, licences, safety, dependability, AI, CAIS, software, railway, space, automotive

This project has received funding from the European Union's Horizon Europe programme under grant agreement number 101069595.

Change Log

Version	Description Change
V0.1	First internal draft for review (EXI)
V0.2	Reviewed document (BSC)
V1.0	Final document (EXI)

Table of Contents

Executive Summary

The Exploitation Report is dedicated to coordinating partners' effort toward the collective and individual exploitation of the project's results. This first version is based on *D.6.2 Exploitation Plan* (made available at M3). This first version presented in September 2023 shall be further extended and updated progressively on an annual basis. In order to be as consistent and as aligned as possible to the indications of the original Exploitation Plan, it basically adopts and extends the original document's structure and content. This allows the Report to be read without consulting the Plan, as it is *de facto* incorporated (and updated) in the Report itself.

The Exploitation Report (based on the initial Plan) analyses the exploitation context and business opportunities to uncover the current and potential market situation. The potential target markets (and target users), as well as the early adopters and followers are identified and analysed, and the competitive environment surrounding the project are assessed. Factors that may influence the exploitation of the results (such as the Technology Readiness Level (TRL), integration, standardization, regulatory aspects, licensing, etc.) are identified and monitored. This iterative work identifies and consolidates business opportunities, considering both the domains and results where exploitation can start in the short term.

The Exploitation Report (based on the initial Plan) defines a methodology and strategy for the appropriate management of the knowledge generated by the project (IPs) and it monitors and iterates it towards the Exploitation Plan. To this end, this task also aims to elaborate on the joint and individual exploitation plans, based on the exploitation context analysis, and on the identification of the exploitable project assets and results.

1. Exploitation Strategy

Exploitation from a scientific and industrial perspective is of paramount importance for SAFEXPLAIN. The SAFEXPLAIN consortium offers a well-balanced and well-complemented combination of industrial and academic partners that will act as a powerful enabler for prolific exploitation. Exploitation is also fostered through the virtual events and workshop that bring together other related projects and that occur where relevant industrials will attend, hence acting as a form of ample and diverse advisory board.

SAFEXPLAIN identifies the following exploitation channels and activities to maximize exploitation opportunities:

- Identification of **project exploitable assets** as critical activity for the exploitation and the sustainability of the project. Exploitable assets include interim and final results, various evaluation activities and lessons learnt from investigations on Deep Learning (DL) specification, implementation, and Functional Safety (FUSA)-DL interaction, as well as potential business models and exploitation pathways.

- Identification of the **main exploitation routes for the consortium** as a whole, for specific groups of partners sharing similar interests / orientation as well as for each partner individually.

- The **procedures to protect IPR** issues of novel tools and technologies, as well as the integration of preexisting individual technologies when integrated into the SAFEXPLAIN solution, including an accurate analysis of the potential conflict among the different licenses that will coexist (e.g., open-sources vs. proprietary, among the multiple open-source licenses).

- Identification and analysis of **the target users** (early adopters and followers) that may benefit from the project's findings and achievements. This analysis is being done in collaboration with the dissemination task that is already in charge of identifying potential target users of the project outcomes. In the exploitation task, the emphasis is closer to "business development".

- Analysis of the **exploitation context and business opportunities** in application domains in order to consolidate the view on the actual market trend. Although this study considers all potential industrial domains, a direct focus is given to the domains in which SAFEXPLAIN industrial partners have direct business opportunities, i.e. AIKO, EXI, NAV, and IKR. The exploitation activities entail the evaluation of project achievements' acceptability by the business world, by addressing: (1) IPR management, (2) open-source communities for project promotion, (3) the definition of joint exploitation agreements, etc.

- Assessing the **competitive environment surrounding the project,** such as technology readiness, integration, standardization and regulatory, and policy framework at the targeted markets as well as future trends at both social, business and policy level. In

particular, standardization and regulatory aspects are paramount concerns in SAFEXPLAIN and are thus explicitly addressed in the project.

- Development of a **sustainability plan of results** that will offer a path beyond the finalization of this project to exploit the results and open new ways to continue the work. The plan addresses (1) IPR management, (2) open- source communities for project promotion, (3) the definition of joint exploitation agreements, (4) a strategy to influence standards, etc. Sustainability is minimally addressed in this first year of project and shall be addressed starting from the second year.

These activities are already directly included in the SAFEXPLAIN *D6.2 Exploitation Plan* (or shall be in the next versions). This Report was made available at month 3 with the primary objective of allowing for fast feedback on exploitable assets and business opportunities. The exploitation report will be updated along with the exploitation reports at months 24 and 36. The report at month 36 will examine and assess the status of the project final results' exploitation and commercialization, taking into account latest technological evolutions and market changes during project's lifetime. It will also include relevant information from the case studies' results.

The feedback gathered during consortium interaction and discussions with experts in the cross-project events, from other key external stakeholders and experts, and industrial actors and decision makers in the targeted markets will be crucial for addressing strengths (benefits), weaknesses (drawbacks and prerequisites), opportunities (existing conditions suitable for promoting the wide adoption of results) and threats. These considerations constitute the basis for planning for successful exploitation and leads to the identification of mechanisms to achieve the actual widespread adoption of project results.

For this latter point, three partnerships of excellence have been established:

1) the planned **Industrial Advisory Board** has been established and counts with the confirmed participation of senior members from academia and industry from all Europe. These members participate as individuals, while retaining their prestigious and highly representative affiliations from top institutional stakeholders. The first meeting of the IAB is expected to take place in autumn 2023;

2) SAFEXPLAIN has established a Memorandum of Understanding (MOU) with **VDA-QMC** (the German Association of the Automotive Industry e.V. (VDA) dedicated to the development of methods and systems for the automotive industry) based on the common interest in Quality, Safety and Security compliance of novel AI-based solutions. This MOU has allowed the two parties to exchange drafts of public documents before their official publication;

3) The partnership mentioned in 2) has been extended to **intacs** (an international Certification Body for ISO/IEC 15504, ISO/IEC 33000 (SPICE) standards) to collaborate on educational syllabi and joint dissemination activities for the brand new ML/DL model included in ASPICE 4.0.

2. Exploitation Results

Partners already conducted a preliminary analysis of their exploitation and IPR strategy, and it has now been further updated. Table 1, ***Expected Exploitable Technological Items,*** identifies SAFEXPLAIN exploitable technological items that are expected to be produced in the course of the project. For each technological item, there is an initial identification of: (1) the item; (2) the owner; and (3) the license, i.e. open-source or proprietary.

Table 1 Expected Exploitable Technological Items.

Item	Owner	License
FUSA-aware DL libraries and extensions	BSC, IKR, RISE, EXI, NAV, AIKO	Open source (MIT, Apache), proprietary
Development and deployment guidelines for safe AI solutions	IKR, NAV	Proprietary
Research prototype to support verification and validation (V&V) of safety critical CPS embedding DL-based components.	RISE	Open source (MIT)
Performance analysis tools for DL software	BSC	Open source (MIT, Apache)
Low-level library for observability and controllability of the target hardware	BSC	Proprietary
Repository of Explainable AI reference architecture and methods to used for V&V of safety critical applications.	RISE	Open source (MIT, Apache)
Integration interface of DL libraries with FUSA analyses toolset	EXI	Proprietary

The project will produce further results that are relevant for the exploitation strategy:

- **Recommendations for safely deploying DL software solutions in Critical AI-based Systems (CAIS) in automotive, railway and space domains**. Recommendations will cover: (a) Techniques to be applied in different stages of the V cycle (e.g. testing), (b) DL techniques and methods (e.g. specification) based on FUSA assessments, and (c) statistical predictability approaches, hardware observability and configuration guidelines for heterogeneous platform complexity. EXI will bring forward those recommendations in safety standard committees through its experts and will incorporate them in their syllabi as part of the hundreds of courses EXI gives worldwide.

- **Assessment in automotive, railway and space domains**. SAFEXPLAIN will share several technical contributions with internal experts and external certification authorities and certification experts in space, railway and automotive. They will make an assessment of those contributions against specific safety standards (e.g. IEC 61508 / ISO 26262 / EN 5012x, ECSS standards in space). Their review will be a valuable asset providing evidence of the

feasibility of the SAFEXPLAIN safety pattern approach and the FUSA techniques to be used in different stages of the life-cycle. EXI will act as internal expert (EXI is at par with TÜVR, TÜVS), whereas external experts will include certification experts (e.g., TÜVS Rheinland) for auto/rail, and ESA certification experts for space. The former will be subcontracted whereas the latter will be approached by AIKO through their regular interactions.

- **Results from case studies**. The adoption of the SAFEXPLAIN technology requires references and success stories in each CAIS application domain. SAFEXPLAIN will collect and deliver evidence from rail, space, and auto case studies to that end. Especially, the explainability, robustness and traceability properties will be evaluated against all the case studies and promoted by the respective case study partners (NAV for auto, AIKO for space, IKR for railway). Moreover, EXI has access to hundreds of key players in most industrial sectors (especially in the automotive domain) and will also promote and disseminate case study results and technologies with those players.

The project results will be exploited by each partner according to its core objectives (business, societal or academic). By including the key know-how in public deliverables and publications, and providing key technological items as open source, individual and joint exploitation can be carried out by interested partners without mutual dependencies that would otherwise preclude it.

3. IPR Management

SAFEXPLAIN generates research, measurements and engineering data obtained from the system simulations, trials, prototyping and the use of testbeds and labs. SAFEXPLAIN partners are committed to making research data accessible and keeping data F.A.I.R. (Findable, Accessible, Interoperable and Re-usable).

To explain how to access the data, additional supporting documentation will be created. WP6 includes a task (T6.2) where knowledge and IPR management is generated and managed. This task guides the participants on how the results will be identified, reported, and protected from early disclosure, and will ensure that the IPR and data management strategies are well defined and coherently executed.

For this purpose, the *D7.2 Data Management Plan* (DMP) was defined at the beginning of the project, and it will be updated as the project evolves. This is part of the SAFEXPLAIN Management Plan.

- **Access rights to Background knowledge.** To ensure a smooth project execution, the project partners will grant each other and their affiliated companies, royalty-free access rights to their Background and Results for the execution of the project. This will allow the researchers to execute the project to the best of their ability, without being hindered by administrative issues. Access rights to this knowledge will be available to all partners only if they are valuable or useful for carrying out project activities. Information may include (among others) the set of tools, hardware designs and software components integrated in the SAFEXPLAIN architecture. The Consortium Agreement (CA) defines further details concerning the Access Rights for Exploitation to Background and Results.

- **Foreground knowledge and IP ownership.** Results shall be owned by the project partner carrying out the work leading to such Results, independently of whether they can be protected or not. If any Results are created jointly by at least two project partners and it is not possible to distinguish between the contributions of each of the project partners, such Results, including inventions and all related patent applications and patents, will be jointly owned by the contributing project partners. Each partner may use the results and material produced within the project for project purposes provided that such use does not come into conflict with the terms of the project *Grant Agreement* or European legislation. To enhance exploitation of the Consortium Results, each contributing party shall have its own full freedom of action to exploit the joint IP as it wishes, and further the goals of the consortium. To promote this effort, the contributing party will have its own full consideration regarding their use of such joint Results and will be able to exploit the joint IP without the need to account in any way to the other joint contributor(s). Further details concerning jointly owned Results, joint inventions and joint patent applications are addressed in the CA.

- **Transfer of Results.** As results are owned by the project partner carrying out the work leading to such results, each project partner shall have the right to transfer Results to their European affiliated companies without prior notification to the other project partners, while always

protecting and assuring the Access Rights of the other project partners. Such use of results will encourage competitiveness of the EU market by creating broader uses of the results and by opening up markets for the Consortium's Results.

- **Patents.** In case a partner wants to submit a patent application, it will inform the other partners following the process described in the CA. Any conflicts will be addressed following a conflict resolution process described in the CA. Information of patent applications will be made available to the EU through regular management reports. The costs of the patent applications will be covered by the submitters.

- **Software/hardware accessories.** The software and hardware accessories (e.g., tools, components, devices, programs) required by other partners to fulfil the project objectives shall only be used for the purpose of the project. Software products shall be made available free of charge, unless it is a commercially available product, and hardware products at base costs including handling fees and depreciation. All these items shall be deleted or returned after the end of the project. These agreements shall be extended beyond the duration of project only at the discretion of the partner owning the software and hardware accessories.

3.1 Exploitation canvas

Table 2 Exploitation canvas

Specific needs	Expected results	Diss/Exp/Comm measures
• FUSA-aware DL-based solutions needed for CAIS (e.g. autonomous cars) • DL-aware FUSA solutions needed to enable certification of DL-based CAIS • Explainability and traceability needed in DL to make DL FUSA compliant • DL software execution on high- performance platform must be time predictable • Industrial viability must be proven in toolsets (for automation) and case studies (for end user acceptance)	• FUSA-aware DL libraries • FUSA patterns to use DL-based solutions in CAIS • Recommendations for FUSA standards to certify DL software • Tools for DL software verification, and analysis of semantics and internals • Performance analysis tools and libraries for DL software on CAIS • Integration of DL libraries with FUSA analyses toolset • Evidence from case studies	• Proprietary commercial solutions for industry • Open-source technological items to ease contributions to increase TRL • Cross-contamination with AI, Data and Robotics projects and partnerships through frequent virtual meetings and a joint workshop • Participation in certification bodies to push SAFEXPLAIN guidelines • Scientific publications and event participation in all relevant communities (AI, FUSA, CAIS)

		• Demos based on case studies at industrial events
Target groups	**Outcomes**	**Impacts**
• End Users in CAIS, e.g. integrators and OEMs • Technology providers in CAIS, e.g. HW and SW providers, AI software companies and developers • Certification authorities and experts, CAIS/AI research community, Policy makers • General public	• Incorporation of SAFEXPLAIN safety guidelines into certification process • Use of SAFEXPLAIN DL libraries, components and API to develop safety-critical software in CAIS • Higher trust in DL-based solutions for FUSA related systems • Contribute to (Strategic Research, Innovation and Deployment Agenda) SRIDA's "safety- by-design" approach among other SRIDA's objectives	• European industry enables fully-autonomous CAIS (e.g. cars, trains, satellites) with certified and economically viable solutions • Increased efficiency of CAIS systems due to safe DL solutions reduces CO2 emissions (up to 80% for different types of vehicles) • European CAIS benefit from DL functionalities and remain competitive in future, while still being trustable

4. Unique Selling Propositions (USPs)

With over a million copies sold, Geoffrey Moore's *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers* (1991, revised 1999 and 2014) is still one of the must-read books for technology marketing leaders. In it, Moore examines the market dynamics faced by innovative new products, with a particular focus on the "chasm" or adoption gap that lies between early and mainstream markets.

The book offers decision-making guidelines for investors, engineers, enterprise executives, marketers and managers throughout the high-tech community. Real-world examples of companies that have struggled in the chasm are also provided.

The core of the book has always been its simple but effective framework for establishing a compelling and clearly differentiated Unique Selling Propositions (USPs).

Here's Moore's original 6-step ***Unique Selling Proposition*** template:

- For (*target customer*)
- Who (*statement of need or opportunity*)
- The (*product name*) is a (*product category*)
- That (*statement of key benefit* - that is, *compelling reason to buy*)
- Unlike (*primary competitive alternative*)
- Our product (*statement of primary differentiation*).

Here's the original ***Unique Selling Proposition*** that Moore made for *Silicon Graphics Inc.*:

- For movie producers and others
- Who depend heavily on post-production special effects,
- Silicon Graphics provides computer workstations
- That integrate digital fantasies with actual film footage.
- Unlike any other vendor of computer workstations,
- SGI has made a no-compromise commitment to meeting film-makers' post-production needs.

Many of the critical elements indicated above in the Exploitation Strategy are captured in this simple, elegant, effective formula. The SAFEXPLAIN consortium has therefore decided to adopt this evergreen model to describe the Unique Selling Propositions (USPs) of its exploitable results.

In the first edition of the Exploitation Plan, the consortium was already able to present **12 USPs**. All partners created at least one USP, some of them up to three. In this first edition of the Exploitation Report, the overall number is **11 USPs.** This number is due to the major evolution and consolidation of the initial USPs: all of them have been revised (a few with minor editing, most with major rehauling, including the change of status for two of them from *individual partner* USPs to *bilateral* USPs).

4.1 EXI01

Item ID	EXI01
Partner	EXI
For	embedded software developers in automotive, railway and aerospace sectors,
Who	need to use advanced ML and DL techniques for highly-dependable systems,
The	**SIL-AI** is a module of the SILcal tool family
That	provides objective evidence to seek compliance for functional safety standards like ISO 26262, ISO 21448 (SOTIF), IEC 61508 and others.
Unlike	the traditional safety analysis tools like *Medini, APIS, SOX, ITEM* and the basic SILCal tool iself,
Our	SIL-AI offers new features and techniques for Verification & Validation uniquely suitable to novel ML/DL-based SW solutions.

Exploitation review findings:

Meetings dedicated to the revision and improvement of this USP were held on 29/03/23, 30/06/23, 01/09/23.

The exact relationship between the existing SILCal X (background IP) and the new SIL-AI (foreground IP) has been thoroughly discussed and a slightly different presentation is now included in the USP.

The competition analysis has been revised and it now has a slightly different priority (Medini currently seems to be most established competitor) and a new identified competitor (ITEM). However, positioning against them seems to be unaffected. In order to make the USP clearer - and more homogeneous with other USPs (see below) - the existing SILCal tool itself is mentioned as competitor.

The statement of primary differentiation has been made more specific, however, this is an area where more elaboration is expected in the next period.

4.2 BSC01

Item ID	BSC01
Partner	BSC
For	embedded software developers and V&V (Verification & Validation) engineers in automotive, railway and aerospace sectors
Who	need to use configure and collect information on hardware events on the NVIDIA® Jetson Orin™
The	**Orin-PMULib** is a dedicated Performance Montoring Unit Library
That	allows configuration on target performance monitoring counters and debug devices.
Unlike	the generic and high-level performance monitoring library solutions like *perf, oprofile, perfmon2, or PAPI*
Our	Orin-PMULib is specifically adapted to the platform and provides a lightweight but accurate way to configure and retrieve precise information on traceable hardware events.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 05/07/23.

This USP version already incorporates adjustments based on early findings that were indicated in the Exploitation Plan, and it has been remarkably stable since its first definition.

Competition analysis is quite complete and initial thoughts on a licencing model have already been discussed internally.

4.3 BSC02

Item ID	BSC02
Partner	BSC
For	embedded software developers and V&V (Verification & Validation) engineers in automotive, railway and aerospace sectors
Who	need to characterize the performance of advanced ML and DL solutions for highly-dependable systems,
The	**pWCET-AI** is a novel probabilistic timing analysis tool
That	allows to characterize the timing behaviour and to derive probabilistic Worst-Case Execution Time (pWCET) estimates of AI-based solutions.
Unlike	existing tools based on traditional deterministic timing analysis approaches, such as static timing analysis (e.g. AbsInt aIT), dynamic analysis (e.g. RapiTime, SymTA/S, AbsInt Timeweaver) or exploiting existing probabilistic methods, such as those based on Extreme Value Theory (e.g. MBPTA-CV, RocqStat)
Our	pWCET-AI allows for trustworthy and tight execution time bounds capturing the specific non-deterministic traits of ML and DL software solutions running on complex SoCs such as, for example, the NVIDIA® Jetson Orin™.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 05/07/23.

This USP version already incorporates adjustments based on early findings indicated in the Exploitation Plan, and it has been remarkably stable since its first definition.

The competition analysis was previously non-existent and is now quite complete (even with categories of competitors) and initial thoughts on a licencing model already being discussed internally.

4.4 BSC03

Item ID	BSC03
Partner	BSC and EXI
For	embedded software developers and V&V (Verification & Validation) engineers in automotive, railway and aerospace sectors
Who	need to build explainable and traceable DL components to be integrated in their systems
The	**DLETlib** is a dedicated DL Explainable and Traceable library, incorporating a strongly structured and layered software architectural design
That	allows for the development of DL components following the requirements from functional safety standards like ISO 26262, ISO 21448 (SOTIF), IEC 61508 and others.
Unlike	traditional DL frameworks (e.g. TensorFlow, PyTorch or Caffe) that only focus on creating a DL infrastructure without supporting explainability/traceability features
Our	DLETlib provides an extension to popular AI frameworks (similar to TensorFlow-probability) to accelerate the adoption of safety standards when DL is used.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 05/07/23, 30/08/23.

This USP version already incorporates adjustments based on early findings indicated in the Exploitation Plan.

In July 2023, EXI realized that some of its project results, which were not expected to contribute to its own EXI01 USP, could be integrated with the BSC's exploitable item described in this USP. EXI approached BSC with a proposal for transforming this *individual* USP into a bilateral, collaborative USP between BSC and EXI. Negotiations followed and a final version was agreed upon based on an intended Open Source licensing model.

NOTE: collaborative, multilateral exploitation agreements are always inherently more complicated than individual ones so it is expected that considerable effort will be devoted to sort out the details in the next period.

4.5 AIKO01

This USP has been cancelled. It was present in the Exploitation Plan but AIKO has decided (with the support of the Exploitation Manager) to only focus on the following AIKO02 USP.

(Note: Original AIKO01 description can be retrieved in D6.2)

4.6 AIKO02

Item ID	AIKO02
Partner	AIKO
For	space industry companies employing assets which require navigation and control (Earth Observation, Telecommunications, Space Debris Collection and Removal, In-Orbit Servicing, etc.)
Who	need algorithms for enabling autonomy in their missions
The	JANE autonomous navigation application is an AI software implementing algorithms for navigation
That	makes space assets more autonomous and reactive and reduces the effort of ground staff
Unlike	AI-solutions for autonomous navigation, pose estimation and object detection developed by innovative companies like SCOUT Space, Rogue Space Systems and LMO Space
Our	AIKO autonomous navigation application enables space critical systems with safe and explainable AI for their navigation operations, complaint to ECSS standards for space software verification and validation (ECSS-E-ST-10-02C), dependability and safety (ECSS-Q-HB-80-03A).

Exploitation review findings:

This USP is a new compared with the previous AIKO02 USP in the Exploitation Plan.

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 28/06/23.

This USP has been deeply revised and now addresses all previously unaddressed findings from the first review.

The beneficiaries section now precisely identifies significant categories.

A tentative name that is better specified has been given to the exploitable result.

Three potential competitors have been identified (none were previously identified).

The statement of primary differentiation has been made more specific now, including the mention of specific sectoral standards.

4.7 RISE01

Item ID	RISE01
Partner	RISE
For	Automotive industry, research communities, testbeds, academia
Who	need to use AI in safety critical applications
The	**XAI for safety** research platform: In the form of an opensource repository (GPL3 license or equivalent) containing methodology and tools enabling the use of XAI techniques to support safety assurance of AI based systems.
That	provides a knowledge base and approach for using XAI to support the application of AI-based components in safety critical systems.
Unlike	Existing best practices and opensource libraries about explainable AI (such as Alibi, AIX360, Xplique) and/or AI safety assurance process (such as SOTIF, AMLAS, ASPICE ML-Model)
Our	**XAI for safety** focuses on a systematic approach to applying Explainable AI techniques for supporting different stages of the DL safety lifecycle (SDSL), including a proposal of evaluation metrics.

Exploitation review findings:

This USP is a new one compared with the previous RISE01 USP in the Exploitation Plan.

Meetings dedicated to the revision and improvement of this USP were held on 29/03/2023, 23/06/2023.

This USP has been deeply revised and addresses all previously unaddressed findings from the first review.

The beneficiaries section now has expanded categories.

The exploitable result, although maintaining its generic category of "research platform", it's further specified as an open source repository.

Following a long discussion on the meaning of "competition" in a purely academic context, six "competitors" have been identified, of which three are Open Sources AI libraries:

- Alibi (https://github.com/SeldonIO/alibi)
- AIX360 (https://github.com/Trusted-AI/AIX360)
- Xplique (https://github.com/deel-ai/xplique)

and three are " safety assurance" standards/models.

The statement of primary differentiation has been expanded and made more specific. In the next period it is expected to verify if and how this USP actually includes all RISE's exploitable results.

4.8 IKR01

Item ID	IKR01
Partner	IKERLAN
For	dependable and Critical autonomous AI-Based Systems (CAIS) developers in the automotive, railway, industrial and aerospace sectors
Who	need to develop and safety certify automated, heteronomous or autonomous systems integrating DL components
The	**Safety Pattern Library (SPL)** is a basic technical reference foundation that provides a set of documented exemplary safety-case(s) and exemplary safety-concept(s), with a technical focus on safety and XAI
That	describe common safety design approaches (solutions) to common design requirements (recurrent problems).
Unlike	the current need to define system-specific designs and argumentations (from scratch) due to a lack of formalized (public) 'reference foundations' and lack of mature safety standards (e.g., ISO 5469 draft)
Our	SPL provides a basic set of documented 'FUSA patterns to use in DL-based solutions', with a subset of them assessed by internal/external experts (e.g., TÜVR) as part of the safety-case assessment(s) (e.g., railway). **SPL** is complementary to **SDSL**.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 29/06/23.

This USP version already incorporates adjustments based on early findings indicated in the Exploitation Plan, and it has been remarkably stable since its first definition.

The competition analysis is now more articulated but still only shows 'abstract' competition, rather than specific competitors. Given the peculiar nature of this exploitable result, this is largely understandable, but it is still a weakness to be addressed in the next period.

4.9 IKR02

Item ID	IKR02
Partner	IKERLAN
For	dependable and Critical autonomous AI-Based Systems (CAIS) developers in the automotive, railway, industrial and aerospace sectors
Who	need to develop and certify for safety: automated, heteronomous or autonomous systems integrating DL components
The	**Safety Lifecycle for DL-software (SLDL)** development is a 'safety lifecycle' (procedures, guidelines, templates) defined in compliance with existing AI-safety drafts (e.g., ISO 5469)
That	provides the required basic procedures, guidelines and templates to support the development of DL-components for CAIS systems, with a technical focus on safety and XAI.
Unlike	the lifecycles in Functional Safety Managements (FSM) for FUSA standards, that do not explicitly consider DL-software
Our	SLDL provides a starting point for developing DL-based dependable CAIS systems, which can be integrated as an extension to traditional Functional Safety Management (FSM) (e.g., IEC 61508). **SDSL** is complementary to **SPL**.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/2023, 29/06/2023.

This USP version already incorporates adjustments based on early findings indicated in the Exploitation Plan, and it has been remarkably stable since its first definition.

Competition analysis is now more articulated but still only shows 'abstract' competition, rather than specific competitors. Given the peculiar nature of this exploitable result, this is largely understandable, but it is still a weakness to be addressed in the next period.

The relationship between IKR01 and IKR02 has to be further clarified. The statement of "complementarity" is OK but can only be accepted in intermediate versions, as final USPs are expected to be self-contained.

4.10 IKR03

Item ID	IKR03
Partner	IKERLAN and EXI
For	dependable and Critical autonomous AI-Based Systems (CAIS) developers in the automotive, railway, industrial and aerospace sectors
Who	need to develop and certify safety automated, heteronomous or autonomous systems integrating DL components
The	**Safety YOLO** library is a basic software re-design/implementation of a subset of YOLO functions in compliance with FUSA standards requirements against SW systematic errors
That	provides a safety software implementation of a subset of YOLO functions for the safe execution of DL-models.
Unlike	software implementations of DL-libraries such as the basic YOLO library itself
Our	safety YOLO library provides a safety software design and implementation, that integrates a structured and layered software architecture, for the deployment of DL-components.

Exploitation review findings:

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 29/06/23, 30/08/23.

This USP version already incorporates adjustments based on early findings indicated in the Exploitation Plan.

In July, EXI realized that some of its project results that were not expected to contribute to its own EXI01 USP could be integrated with IKR's exploitable item described in this USP. EXI approached IKR with a proposal to transform this *individual* USP into a bilateral, collaborative USP between BSC and IKR. Negotiations followed and a final version was agreed upon based on an intended Open Source licensing model.

NOTE: collaborative, multilateral exploitation agreements are always inherently more complicated than individual ones so it is expected that considerable effort will be devoted to sort out the details in the next period.

4.11 NAV01

Item ID	NAV01
Partner	NAVINFO
For	Critical autonomous AI-Based Systems (CAIS) solution providers (Engineers, AI researchers and developers) for the autonomous industry
Who	plan to design vision-based models with explainability and safety in mind and evaluate the reliability and safety of their AI system
The	**NIE GuardAI** is a platform comprised of specialized AI modules for developing and assessing safe and explainable ADAS systems that complies to the FUSA and explainability requirements.
That	provides the foundational AI components for developing explainable perception models (e.g., CNN layer, pretrained feature extractors, decoders), and evaluation protocols for extensively assessing them for adherence to functional safety.
Unlike	Other software tools (e.g. PyTorch), which utilizes DL components that do not adhere to the safety and evaluation frameworks (e.g Neurocat, Bosch AI Shield) that only assess the models for adversarial robustness under laboratory settings.
Our	Framework allows the solution providers to develop and evaluate customized production-ready AI based perception systems for their specific use cases and conduct exhaustive safety and explainability assessment on a wider range of vision task in real-world setting, allowing the overall AI system to be safety certified. Additionally, it provides easy integration in existing MLop workflows.

Exploitation review findings:

This USP is a new one that was not in the previous NAV01 USP in Exploitation Plan.

Meetings dedicated to the revision and improvement of this USP were held on 29/03/23, 02/06/23.

This USP has been deeply revised and addresses most of previously unaddressed findings from the first review.

The beneficiaries section now identifies significant categories more precisely.

A new tentative name has been given to the exploitable result that is better specified both as positioning and as characterization.

Three potential competitors in two categories have been identified (none had been identified previously).

The statement of primary differentiation has been made more elaborated.

At end of August 2023, NAV announced a major internal organizational change, potentially impacting this USP: a meeting with the new NAV responsible is expected to be organized at M13.

4.12 NAV02

Item ID	NAV02
Partner	NAVINFO
For	Critical autonomous AI-Based Systems (CAIS) solution providers for autonomous driving industry
Who	Who want to ensure the safety and reliability of their autonomous driving system and achieve relevant safety certification
The	**NIE Safe AI development and deployment guidelines** on AI based models and the hardware and software infrastructure provides requirements and relevant quantitaive and qualitative metrics on computer vision-based detection tasks for autonomous driving systems.
That	provides guidance and recommendations on the development and deployment of DNN-based AI perception models that comply to the safety, traceability and explainability requirements and provides the key performance indicators and performance metrics (including uncertainty estimation, failure detection, robustness, fault tolerence, etc).
Unlike	the other available AI solutions and services offered by AI-based vision service companies like MobilEye Supervision that lack compliance with explainabiliity and tracelability requirements for safety certifications and deployment of AI systems
Our	**NIE Safe AI development and deployment** allows users to specify and customize automobile use case and get information and recommendations that best suit their specific use-case throughout the AI development and deployment life cycle.

Exploitation review findings:

This USP is a new one, if compared with the previous NAV02 USP in Exploitation Plan.

Meetings dedicated to revision and improvement of this USP were held on 29/03/23, 02/06/23.

This USP has been deeply revised and addresses most of previously unaddressed findings from the first review.

The beneficiaries section now identifies significant categories more precisely.

A new tentative name has been given to the exploitable result that is better specified both as positioning and as characterization.

One potential competitor has been identified (none has been identified previously), although in quite generic terms.

The statement of primary differentiation has been elaborated on.

At end of August 2023, NAV announced a major internal organizational change, potentially impacting this USP: a meeting with the new NAV responsible is expected to be organized at M13.

5. Beyond initial Unique Selling Propositions

In the nine months between the release of the Exploitation Plan and of the Initial Exploitation Report, the partners have mostly focused on the updated/extended/upgraded Unique Selling Propositions, in light of the first internal reviews and of the first results of the project itself.

Partners are aware that Moore's template is a very smart starting point for elaborating on exploitation strategies but that it does not cover all aspects that might be relevant for a full exploitation of results. The above mentioned 'exploitation canvas' actually covers a wider spectrum of concepts.

During the kick off meeting, besides and beyond Moore's template, a decalogue encompassing all exploitation aspects was presented:

#	*EXPLOITATION ASSETS*
1	**a catalogue of** (foreground/background) **Intellectual Properties** (associated with the specific exploitable IP)
2	**their origin** (developed internally, by partners, in collaborative projects, etc...)
3	**their ownership** (fully owned, jointly owned, public domain, etc…)
4	**their licensing status** (free, nominal, discounted, premium; all/some rights reserved; geographical scope, etc...)
5	**their value** (by development cost, internal estimation, external audit, price list...)
6	**their protection** (copyrights, trademarks, patents, NDAs, etc…)
7	**policy on physically** (embedded in actual physical goods) **and virtually distributed** (internet or mobile platforms) **IPs**
8	**policy on returns from products** (physical units), **services** (subscriptions, uses, resources...) and **licensing** (upfront fee, recurrent fees) **of IPs**
9	**policy on the so-called** "some rights reserved" , or "**open content**" **approach** (open software, creative commons...)
10	**policy on the so-called "freemium" biz-model (**giving away some/all of your IPs for free, to open a market and/or to build on revenue from other associated services)

While not necessarily all these assets will have to be defined in detail for all the exploitable results, a consideration to them will be given and succinct descriptions (including justifications for exclusion) will be collected, to complement the enhanced Unique Selling Propositions. It is also important to note that while USPs are thought of as marketing tools to be fully externally exposed and therefore highly *public*, some of these assets might well be *restricted* or *confidential*.

In the period between the release of the Exploitation Plan and of the Initial Exploitation Report, given the nature of the identified exploitable items themselves (basically structured texts and code

libraries, hence *immaterial, soft IPs* to be essentially protected by copyrights and licences), most of the extra effort has gone to LICENCING and related aspects (therefore, esp. 1, 3, 4, 6, 9).

6. Key performance indicators (KPIs)

As well known by seasoned experts as well as by experienced practitioners, the exploitation success is extremely difficult to predict and monitor. Control of it is only minimal within the project and externalities are overwhelming.

The consortium proposed a set of KPI in the Exploitation Plan, which are directly linked to what was exposed in the Exploitation Plan itself. We now confirm the original list. The first set of tentative metrics is presented in this Initial Exploitation Report at M12, based on estimation and evaluation carried out in the first year of life of the project, and the first quantitative monitoring results will be presented at M24.

KPI Short name	KPI description	Measure
#01 Exploitable Results	identified and named exploitable results	10
#02 Categories of products	impacted categories of products in the market	4
#03 Competitive Products	primary competitive products	30
#04 Competitive Orgs	primary competitive organisations	20
#05 Target User Groups	target groups of customers/users	6
#06 User-case Scenarios	opportunities in user-case scenarios	3
#07 Proprietary Products	proprietary products announcements	5
#08 Open-Source Products	open-source products announcements	5

The above-mentioned KPIs will be carefully monitored and revised yearly, as they may change or evolve based on the project's progress.

Acronyms and Abbreviations

- CA – Consortium Agreement
- D – Deliverable
- DoA – Description of Action (Annex 1 of the Grant Agreement)
- CAIS - Critical AI-based Systems
- EB – Executive Board
- EC – European Commission
- FuSa – Functional Safety
- GA – General Assembly / Grant Agreement
- HPC – High Performance Computing
- IPR – Intellectual Property Right
- KPI – Key Performance Indicator
- M – Month
- MS – Milestones
- PM – Person month / Project manager
- OEM - Original Equipment Manufacturer
- WP – Work Package
- WPL – Work Package Leader

7. ANNEX: A basic primer on Intellectual Property

7.1 Intellectual Property and associated Rights (IP and IPRs)

Intellectual property (IP) is a term referring to a number of distinct types of creations of the mind for which a set of exclusive rights are recognized — and the corresponding fields of law.

Under intellectual property law (internationally acknowledged but with significant national flavours), owners are granted certain exclusive rights to a variety of intangible assets. Common types of intellectual property include copyrights, patents, trademarks.

Intellectual property rights (IPRs) can be bought and sold, leased or rented, or otherwise transferred between parties in much the same way that rights to real property or other personal property can be transferred.

7.2 What are copyrights and what is granted by them?

Copyrights protect products of the mind like: writings, music, sculpture, computer software, graphs, drawings, and mask works.

For a work to be copyrightable, it must be an original work of authorship fixed in any tangible medium of expression.

"Fixed" expressions include a broad range of works, including printouts, computer code, computer chips, and photographs.

Copyrights provide exclusive rights to authors or their assignees.

A copyright is used to prevent others from reproducing, distributing, performing or displaying publicly, or preparing ***derivative works*** without permission of the author.

7.3 What is a patent and what is granted by it? What about SW?

A ***patent*** is an agreement between the government and the inventor whereby, in exchange for the inventor's complete disclosure of the invention, the government gives the inventor the right to exclude others from making, using, or selling the invention.

Can software be patented?

- it is normally protected through copyrights, which do not protect the idea, but the expression of the idea
- certain types of software can be patented (US and EU very different approaches here, UE much more restrictive), and it may be preferable in certain situations for software to receive both copyright and patent protection

- one cannot patent an algorithm to perform mathematical functions or operations in software

7.4 What are trademarks?

Trademarks are any word, name, or symbol, or any combination of these elements (including smells for perfumes, noises for engines…), that are used to identify goods.

Trademarks provide some protection to its owner from those who would attempt to trade on the goodwill and recognition established by use of the same or a similar mark.

Trademarks can be registered or not.

7.5 What is licensing?

Licensing is the act of granting somebody else (some aspects of) the use of own IPR.

All kinds of IPRs (patents, copyrights, trademarks) are subject to licensing.

IPRs are likely to be registered/formalized to be better protected against infringement.

IPRs can be costly (with up front and/or recurrent fee) but can also be free (FLOSS).

Licensing does not affect ownership of IPR.