PROJECT RESULTS

STARLIGHT

Sustainable Autonomy and Resilience for LEAs using AI against High Priority Threats

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No. 101021797.

FOREWORD

Sue O'Sullivan
Chair, **INVICTM**: International Network Supporting Victims of Terrorism and Mass Violence
Member, **STARLIGHT External Advisory Board**

As Artificial Intelligence (AI) becomes increasingly embedded in policing and public safety across Europe and globally, our responsibility is clear: innovation must advance alongside fairness, transparency, and respect for victims' rights. Emerging technologies can strengthen decision-making and improve public safety, but they can also introduce risks that affect victims in profound ways, particularly when systems lack oversight, clarity, or safeguards to protect sensitive information.

For civil society and victim-support organisations, concerns around rights, privacy, bias, data integrity, and the erosion of human judgement are not abstract. For victims of crime, a breach of confidentiality, an unexplained automated decision, or a system that misinterprets their information can retraumatise, discourage reporting, and erode confidence in justice processes. Trust is fragile, and once lost, difficult to rebuild. Yet when AI is developed responsibly, with transparency, strong governance, and meaningful engagement, it has the potential to improve access to support, strengthen investigations, and enhance communication with victims and communities. Ethical innovation can open pathways to better safety and more coordinated services, provided it remains firmly anchored in human rights, reasonableness and proportionality.

The STARLIGHT project serves as a good example of this approach. Open dialogue on the design, implementation, and governance of AI, a dialogue that welcomes scrutiny, fosters multidisciplinary collaboration, and evaluates ethical and legal safeguards, is essential to ensuring that the future of technology in policing is both responsible and rights-based. INVICTM's experience across jurisdictions shows that progress is strongest when victims, survivors, and civil society organisations are engaged from the outset. Their insight highlights unintended harms, strengthens accountability, and ensures that new technologies reflect values of dignity, equality, and compassion.

True progress in public safety is not created for victims, but with them. By embedding their perspectives into the design, governance, and evaluation of AI tools, we build systems that protect rights, reinforce trust, and serve communities with integrity. The results of STARLIGHT provide a substantive foundation for continued discussion on rights-based AI deployment. As the project concludes, its lessons remind us that ethical innovation depends on transparency, partnership, and an unwavering commitment to those most impacted.

I commend members of the STARLIGHT consortium for your dedication, leadership, and invaluable contributions to this important work, and for advancing thoughtful, responsible approaches that can continue to guide future efforts across Europe and beyond.

WELCOME MESSAGE

Dr Nizar Touleimat
STARLIGHT Project Coordinator, **CEA-List**

It is my great pleasure to welcome you to the STARLIGHT results brochure. Over the past 52 months, our consortium of fifty partners has worked together with a single ambition: to strengthen Europe's security ecosystem through responsible, transparent, and human-centred artificial intelligence. This brochure brings together the achievements of that collective effort.

STARLIGHT was launched at a moment of significant transformation for Law Enforcement Agencies (LEAs) across Europe. The rapid growth of data, the complexity of criminal activity, and the increasingly hybrid nature of security threats demanded new capabilities, new methods, and new forms of cooperation. From the beginning, our vision was clear: AI solutions must not only be technically advanced but aligned with European values, built on legal and ethical foundations, and designed to empower, not replace, the practitioners who protect our societies.

A key pillar of STARLIGHT's success has been its collaborative, iterative co-development model. By bringing together LEAs, researchers, and industry partners from the earliest stages of design, we ensured that each AI capability was grounded in operational reality, tested in authentic environments, and refined through direct feedback.

The results presented in this brochure reflect the progress we have made: interoperable operational tools, datasets, governance frameworks, pilot evaluations, and a sustainable community that will continue supporting innovation well beyond the project's completion.

I thank all partners, contributors, and stakeholders for their dedication and professionalism. STARLIGHT's legacy will support a safer, more resilient, and more secure Europe for years to come.

EXECUTIVE SUMMARY

STARLIGHT has united fifty partners from law enforcement, research, industry, and civil society in a shared mission: to advance responsible and trustworthy artificial intelligence for security in Europe. The project was conceived at a time when the scale and complexity of data required new approaches to prevention, detection, investigation, and protection. STARLIGHT responded by building a fully integrated AI ecosystem that spans datasets, tools, governance, operational pilots, and long-term sustainability mechanisms.

Central to STARLIGHT's achievements is its co-development approach, which brought LEAs and technical partners together in iterative development cycles. This ensured that every AI-based tool was informed by real operational needs and tested across six priority security domains: counter-terrorism, child sexual exploitation, border management, cybersecurity and cybercrime, serious and organised crime, and the protection of public spaces.

To safeguard responsible innovation, STARLIGHT embedded transparency, accountability, privacy, and human oversight throughout development. The integration of the AP4AI approach[1] provided a structured pathway for aligning emerging capabilities with the expectations of the EU AI Act and GDPR (General Data Protection Regulation), reinforcing the project's commitment to trustworthy, human-centred design.

Beyond individual tools, STARLIGHT has built a sustainable legacy. The project created a repository of high-quality datasets, an AI Community of Expertise, with its results designed for integration into future-facing infrastructures such as the Europol Tool Repository, the European Anti-Cybercrime Technology Development Association (EACTDA) and Tools4LEAs. These platforms will ensure continued access to STARLIGHT's outcomes, enabling ongoing evaluation, refinement, and adoption.

The results presented in this brochure reflect a coordinated, multidisciplinary effort to ensure that AI can serve as a responsible, effective, and transparent enabler of security in Europe. STARLIGHT leaves behind not only innovations, but a methodology and community that will continue strengthening Europe's strategic autonomy in AI for LEAs.

[1]www.ap4ai.eu

CONTENTS

Safeguarding rights, strengthening trust: victim engagement and collaboration in the age of AI policing

Sue O'SULLIVAN
INVICTM Chair
STARLIGHT
External Advisory Board Member

Levent ALTAN
Victim Support Europe
INVICTM Member

Dr An VERELST
Public Health Belgium
INVICTM Member

Exploring how ethical AI, inclusive partnerships, and victims' voices can guide law enforcement toward safer, fairer, and more trusted public safety outcomes.

Sue O'Sullivan, Chair, INVICTM

As artificial intelligence (AI) becomes increasingly embedded in law enforcement, its potential to enhance safety, decision-making, and operational effectiveness must be balanced with respect for human rights, fairness, and accountability. Ethical safeguards, transparency, and a victim-centred approach can be a bridge between innovation and public trust.

While AI holds promising opportunities to facilitate victims exercising their rights, it can pose a risk to victims accessing justice.

Civil society organisations broadly recognise AI's potential to enhance public safety and operational efficiency, however, raise concerns about how these technologies are developed, deployed, and governed, particularly when used by law enforcement agencies.

These concerns are grounded in fundamental human rights principles and can be summarised through several interrelated themes: protection, privacy, bias, accountability, transparency, and public trust.

From a civil society and victim-support perspective, several key concerns emerge:

- **Privacy and data protection**
- **Bias and discrimination**
- **Lack of transparency**
- **Erosion of human judgement**
- **Effect on reporting.**

The importance of engaging civil society and victim support organisations The strength of INVICTM lies in the ability to share lessons learned, leverage the knowledge and expertise of our members and their networks, to influence change and turn research into action within our own countries and globally.

The opinions, findings, conclusions, and recommendations expressed in this article are those of the contributors and do not necessarily represent the official positions or policies of individual INVICTM members, agencies, governments, or affiliations.

In our work, we have seen how inclusive collaboration strengthens both public trust and operational outcomes, both from a victim-support and a law enforcement perspective.

When victims and communities see that their safety, privacy, and dignity are prioritised, they are more likely to engage with and support law enforcement processes. Law enforcement agencies must involve victims' voices and civil society early in the process to ensure AI tools reflect shared values of fairness and respect for victims' rights.

Active and ongoing engagement with civil society and victim support organisations is essential for several reasons:

- **Building trust:** open communication fosters transparency and confidence in how AI is applied.
- **Co-designing ethical frameworks**: collaborative development ensures safeguards are practical and rights-based**.**
- **Informing policy and training**: civil society insights help shape guidance and operational readiness.
- **Public accountability**: oversight from external partners enhances legitimacy and social acceptance.

Civil society engagement in AI for law enforcement takes many forms. For instance, on 21 May 2025, the STARLIGHT project participated in the INVICTM International Forum in Lisbon, where experts on the panel highlighted legal, ethical, and operational safeguards while discussing the importance of including victims and stakeholders in the use of AI.

Another key example is CENTRIC's landmark UK survey of over 10,000 citizens, *Citizen Consultation on the Police Use of AI in the UK: Acceptance, Trust, and Accountability*, which provides valuable insights into public trust, acceptance, and accountability. Findings underscore where engagement and oversight are most needed to ensure responsible AI use in policing.

Ethical AI begins when those it affects most have a voice in shaping it. This principle reflects a call from civil society: decisions about technology, particularly those influencing safety, justice, and rights, must be made with the participation of those most impacted.

Safeguards for ethical and accountable AI When AI tools in law enforcement operate without transparency, individuals cannot challenge decisions that affect them.

Clear disclosure of AI use, combined with independent oversight, is essential to ensure accountability and protect rights.

Key safeguards for AI in law enforcement include:

- **Transparency and explainability**: clear communication about how systems function and how decisions are reached.
- **Human oversight and accountability**: ensuring that trained professionals remain responsible for final decisions.
- **Legal and ethical frameworks**: embedding compliance with rights-based and proportionality principles.
- **Victim-centred protections**: designing systems that prioritise safety, dignity, and informed consent.
- **Independent review and oversight**: external bodies to assess effectiveness, equity, and rights compliance.
- **Bias auditing and inclusive design**: regular testing to identify and mitigate discriminatory outcomes.
- **Data minimisation and protection**: collecting only what is necessary and ensuring robust data security.

The Ethics, Legal, and Compliance Monitoring and Assessment Methodology developed within the STARLIGHT project provides an example of a structured approach that embeds oversight, transparency, and respect for fundamental rights into every stage of AI adoption, ensuring accountability and ethical integrity.

The positive potential of AI for public safety and victim support Despite the risks, AI also offers tremendous potential when used responsibly and ethically. It can help identify emerging threats, enhance crisis communication, and strengthen coordination between agencies.

In victim support, AI can improve accessibility and enable faster responses, translation, and connection to appropriate support and services. When ethical safeguards are applied, AI technologies hold meaningful potential to:

- **Enhance public safety**: improve situational awareness, early warning systems, and threat detection.
- **Support victims more effectively**: enable faster identification and improved service coordination.
- **Aid investigations**: assist in analysing digital evidence and detecting patterns in complex cases.
- **Improve accessibility**: provide multilingual and adaptive communication tools for victims.
- **Strengthen resilience**: help agencies anticipate needs and plan long-term support strategies.

AI-driven tools, when grounded in transparency and ethics, can strengthen both justice outcomes and public trust, provided they are guided by human judgement, empathy, and respect for rights.

Conclusion Safeguarding rights and strengthening trust in the age of AI policing requires more than technological innovation; it demands meaningful engagement with victims and civil society at every stage. Law enforcement must be guided by a victim-centred approach rooted in respect, dignity, and equality, one that places the needs, rights, choices, safety, and well-being of victims at its core. Victims' voices should help shape the design, implementation, and evaluation of emerging technologies, ensuring that innovation enhances rather than erodes trust.

By listening to and learning from those most affected, agencies can uphold compassion, fairness, and justice while safeguarding the rights of all. True progress in justice and recovery is built with those most affected - not for them. When victims and survivors are meaningfully engaged at every table, their insights strengthen systems, inspire innovation, and are a bridge to public trust.

Continuous dialogue between law enforcement, technical experts, practitioners, oversight bodies, victims, survivors, and civil society is the cornerstone of ethical progress, accountability, and public trust in AI-driven public safety.

STARLIGHT at the 2025 INVICTM International Forum Supporting Victims of Terrorism and Mass Violence

STARLIGHT: AI innovation shaped by operational needs

Dr Nizar TOULEIMAT
Project Coordinator
STARLIGHT

> LEAs must uphold strict standards of transparency, fairness and accountability. AI technologies used in policing must therefore be explainable, legally sound and operationally justified. These commitments shaped STARLIGHT's guiding philosophy: powerful innovation grounded in European values.

A new landscape for European security Across Europe, law enforcement agencies are confronted with a rapidly evolving security landscape. Investigators must process unprecedented volumes of multimodal data, track offenders who operate seamlessly across borders and platforms, and respond to threats that blend physical and digital environments. These challenges highlight the urgent need for AI solutions that are reliable, explainable, and operationally grounded.

At its core, STARLIGHT aimed to strengthen Europe's strategic autonomy in security-related AI by empowering law enforcement agencies (LEAs) with trustworthy, interoperable, and operationally validated AI capabilities.

Understanding real operational pressures From the start, STARLIGHT grounded its work in the operational realities of its LEA partners. Teams from eighteen countries examined the evolving conditions they face in areas such as counter-terrorism, child sexual exploitation, border and external security, cybersecurity and cybercrime, serious and organised crime and public-space protection. Their assessments highlighted a common challenge: the sheer scale, heterogeneity, and pace of data now far exceed the capabilities of traditional investigative processes.

At the same time, LEAs must uphold strict standards of transparency, fairness and accountability. AI technologies used in policing must therefore be explainable, legally sound and operationally justified. These commitments shaped STARLIGHT's guiding philosophy: powerful innovation grounded in European values.

Six priority domains, one shared purpose To ensure relevance, STARLIGHT focused on six high-priority security domains. These domain areas represent the core operational pillars in which LEAs require advanced analytical support.

In counter-terrorism, analysts face the relentless spread of multilingual extremist content and the need to reconstruct events with speed and precision. In child protection, the fight against exploitation involves massive volumes of imagery, conversations and metadata. Effective situational awareness in border and external security operations requires complex multimodal data and multilingual information.

Cybercrime and cyber-threat assessment require rapid detection of anomalies in constantly shifting digital environments. Serious and organised crime investigations are hampered by “infobesity”, where the volume of seized and contextual data overwhelms analysts. And in the protection of public spaces, authorities must rapidly detect anomalies in crowd movement and behaviour.
These operational domains shaped the expectations and design principles of STARLIGHT’s suite of tools.

Linking STARLIGHT’s tools to real operational needs STARLIGHT did not design generic AI features; it created tools tied directly to the six domain areas and the analytical challenges they represent. Across all six domains, STARLIGHT’s tools share a common foundation: ethics by design, explainability, human-centricity and transparent analytical logic, principles embedded directly into their design.

Co-development: AI built with, not for, LEAs A defining feature of STARLIGHT is its co-development methodology. Work progressed in six-month cycles where investigators and technical teams collaborated continuously, testing early prototypes, refining interfaces and adjusting functionalities. This agile process ensured that AI evolved from concrete operational feedback rather than abstract research assumptions. Each cycle concluded with a pilot phase in which LEAs evaluated the emerging functionalities in realistic operational scenarios. This iterative rhythm, design, test, refine, formed the backbone of STARLIGHT’s innovation model.

Pilots: bringing AI into the field Pilot scenarios were developed according to international exercise standards to mirror real investigative challenges. Each scenario described the role of participating LEAs, the data required, the intended workflow, and the outcomes expected. During execution, LEAs assessed whether the tools delivered the promised operational value, providing immediate and targeted feedback. This ensured the tools were not only technically advanced but viable within the constraints, needs and tempo of LEA work.

An integrated framework for trustworthy AI The STARLIGHT framework brings together multilingual and multimodal data processing, advanced AI analytics, adversarial robustness and secure collaborative workspaces. It emphasises explainability and accountability throughout, reflecting the legal and ethical expectations of European policing.

Several tools enhance capabilities in counter-terrorism by enabling the analysis of multilingual extremist content, the correlation of cross-platform signals and the rapid reconstruction of events using multimodal evidence.

Tools supporting child sexual exploitation investigations help identify harmful material, detect grooming patterns in conversations, and link imagery or accounts across platforms and jurisdictions.

Border and external security operations rely on complex multimodal data and multilingual information to support situational awareness.

Cybersecurity and cybercrime tools strengthen the detection of malware behaviour, correlate threat indicators and support automated reasoning that reduces the time needed to recognise and mitigate threats.

For serious and organised crime, the tools tame information overload, enabling analysts to derive meaning from large volumes of textual, visual and geolocated information and to detect weak signals that might otherwise be lost in noise.

Tools that analyse crowd flows and anomalies, supporting real-time situational awareness in transit hubs, stadiums and open-air events.

Tools within this framework can gather, fuse and interpret data from diverse environments, online platforms, sensors, cyber infrastructures, and transform it into actionable insights. The framework also includes capabilities to defend LEA systems against manipulation and cyber-attacks.

A European pathway for sustainable innovation Beyond the tools themselves, STARLIGHT's greatest contribution may be the community it built. By uniting LEAs, researchers, industry and policymakers under a shared human-centric approach, the project has demonstrated that innovation and accountability are not opposing forces, but complementary imperatives. This collaborative foundation strengthens Europe's strategic autonomy and lays the groundwork for long-term cooperation, uptake and evolution of AI in security domains.

A legacy of responsible power STARLIGHT shows that Europe can develop sophisticated AI capabilities without compromising its democratic values or legal obligations. By embedding operational realities, ethical safeguards and human judgement into every stage of development, the project has delivered a pathway for trustworthy, accountable AI in security. Its legacy lies not only in tools, but in a culture, one that recognises that the future of AI in policing is not merely technical, but deeply human.

Building trust in AI: STARLIGHT's path to accountability through AP4AI

Prof Babak AKHGAR OBE
Director
Centre of Excellence in Terrorism, Resilience, Intelligence and Organised Crime Research (CENTRIC)

Prof Marco GERCKE
Director
Cybercrime Research Institute (CRI)

> Compliance is not an audit: it's a continuous process of reflection and improvement.
>
> Dr Marco Gercke, CRI

From the outset, STARLIGHT recognised that trustworthy AI in policing cannot be achieved through technical performance alone. Instead, legal compliance must be addressed alongside system design. This became even more relevant when the EU AI Act was passed during the project.

One of the key challenges for STARLIGHT was to operationalise legal and ethics compliance in a large consortium with several dozen CODEVs. To operationalise legal compliance and provide an actionable framework for the consortium, STARLIGHT adopted the Accountability Principles for AI (AP4AI) framework, co-developed by CENTRIC and Europol, as a central methodological reference. AP4AI was designed to address a core challenge for public authorities: how to harness the potential of AI in policing and security while respecting the obligations that European governments owe to their citizens.

Within STARLIGHT, the framework provided a structured way to reflect on this challenge throughout the CODEVs' development process. It also supported assessment against the requirements of the EU AI Act. Rather than treating compliance as a final validation step, AP4AI thus enabled developers to carry out ongoing self-assessments.

The AP4AI framework rests on two complementary foundations. First, its principles were developed by multidisciplinary experts spanning law, technology, human rights, and policing. Second, these principles were subjected to international public consultation, ensuring democratic input beyond the project itself. This combination of expert knowledge and public validation gave AP4AI both technical credibility and social legitimacy.

AP4AI articulates clear expectations for accountable AI in policing, including transparency, explainability, meaningful human oversight, continuous learning, and mechanisms to address errors or unintended harm. In STARLIGHT, these principles informed the project's responsible innovation approach and provided a shared reference point for technical teams, legal experts, and law enforcement partners.

AP4AI in STARLIGHT: turning principles into practice STARLIGHT adopted AP4AI after the introduction of the EU AI Act as its internal accountability-by-design methodology. Instead of waiting for final regulatory guidance, CODEVs used the framework to translate the abstract provisions of the EU AI Act into concrete questions within each cycle. This approach allowed potential compliance issues to be identified and discussed while design choices were still open.

Implementation followed a two-stage process. First, CENTRIC adapted the generic AP4AI framework to the technical and operational context of STARLIGHT. Second, each CODEV completed a structured self-assessment covering risk identification, safeguards, data governance, human oversight, and alignment with anticipated AI Act obligations. These assessments were then discussed in dedicated "sparring rounds" involving the CODEV developers, legal experts from the Cybercrime Research Institute (CRI), and the CENTRIC team.

Through this iterative exchange, legal and accountability considerations became embedded in the technical and operational design of each STARLIGHT tool. AP4AI does not certify legal compliance, nor does it function as a simple checklist. Instead, it provides a systematic process allowing CODEVs to self-assess whether the AI-based systems they develop are compliant with the AI Act.

By September 2024, completion of the AP4AI process became a formal requirement for all tools entering STARLIGHT's operational pilots. This created a transparent and auditable body of documentation that can support future evaluation, deployment decisions, and external scrutiny.

> Through AP4AI, STARLIGHT turned accountability from an abstract principle into a practical, iterative process.
>
> Prof Babak Akhgar OBE, CENTRIC

A legacy for future security research STARLIGHT's integration of AP4AI offers an important lesson for future EU security research projects: embedding accountability from the beginning strengthens governance while also supporting innovation. The framework created a shared space in which developers, law enforcement practitioners, and legal experts could engage in informed, critical discussion about risks, safeguards, and societal expectations.

While responsibility for compliance ultimately remains with individual partners, STARLIGHT provided a common methodology, shared tools, and structured documentation to support that responsibility. Perhaps more importantly, it fostered a culture of continuous reflection, something that can be challenging in technically driven research projects.

By turning accountability from an abstract principle into a practical, iterative process, STARLIGHT makes a concrete contribution to Europe's ongoing debate on trustworthy AI in policing. It demonstrates that responsible AI is a feasible and measurable practice grounded in democratic values and supported by rigorous technical work.

Building ethical, realistic, and compliant data for AI in security

Dr Henri BOUMA
Senior Scientist
TNO

Ezgi EREN
Doctoral Researcher
KU Leuven Centre for IT & IP Law

> Achieving this vision required more than collecting data. It meant constructing an entire ecosystem where datasets could be created, shared, and reused safely, transparently, and in alignment with GDPR and the EU AI Act.
>
> Dr Henri Bouma, TNO

Responsible data: the cornerstone of trustworthy AI From the outset of STARLIGHT, the project's dataset development efforts focused on one of the initiative's most complex challenges: creating realistic and representative datasets for training, testing, and benchmarking Artificial Intelligence (AI) tools in a legally, ethically, and socially compliant manner.

Achieving this goal required not only the technical precision involved in the creation and curation of data but also the establishment of a secure, transparent, and responsible data ecosystem capable of supporting advanced AI development without compromising privacy, rights, or trust.

The consortium approached this mission by focusing on three interconnected objectives: ensuring realistic and representative data, protecting sensitive information through anonymisation or synthetic generation, and preventing bias in dataset creation.

These objectives formed the foundation for a comprehensive strategy combining data collection, generation, and rigorous balancing.

From concept to compliant practice Guided by continuous feedback from Law Enforcement Agencies (LEAs), the team created datasets that reflect realistic operational conditions while maintaining full compliance with ethical and legal standards. Privacy-preserving mechanisms were embedded throughout the process, including anonymisation, synthetic data generation, and secure handling protocols. These measures allowed AI tools to be developed and tested responsibly without exposing personal data.

In total, 16 tools were developed within this effort: five for data collection and annotation, three for anonymisation, and eight for synthetic data generation. Thirteen of these are available in the STARLIGHT repository for the whole consortium, while three can be made available to European LEAs upon request.

This work also included the assessment of more than 200 datasets for legal and ethical aspects. Ninety datasets were uploaded to the repository, including 58 existing or public datasets and 32 new ones, covering collected, anonymised, and generated data.

No dataset was used within STARLIGHT without first undergoing a thorough legal and ethical assessment procedure. Each dataset was documented, reviewed for Ethical, Legal and Social Aspects (ELSA) by partner KU Leuven, and approved following data protection best practices.

This process ensured that every dataset entering or leaving the project complied fully with the General Data Protection Regulation (GDPR) and STARLIGHT's internal ethics framework.

Lessons for future research The project highlighted a key lesson for future R&D: arranging safeguards and approval for a large initiative aiming to create many datasets and AI applications for multiple use cases becomes disproportionately complex compared to the effort for research and development. For this reason, future AI research and development projects focusing on a single application with one clear purpose may ensure that ELSA governance remains more manageable.

This hands-on experience also proved invaluable in resolving crucial mismatches between engineering and legal understandings of personal data, particularly concerning public datasets and the limits of data processing.

Ultimately, this effort delivered a robust data ecosystem that supports innovation, ensures compliance, and contributes directly to STARLIGHT's mission of enabling transparent, safe, and trustworthy AI for European law enforcement.

16

Data-related tools developed

For data collection, anonymisation, and synthetic data generation, delivering practical solutions to create, protect, and enhance data for AI development.

200+

Datasets evaluated

Reviewed and validated for ethical and legal compliance, ensuring every dataset meets GDPR and STARLIGHT's ethics framework before use.

90

Datasets published

Shared through the STARLIGHT repository for LEA research and testing, providing secure, accessible resources to support responsible AI innovation.

When technology meets operational reality

Michalis LAZARIDIS
Researcher
Centre for Research and Technology Hellas (CERTH)

Filipe APOLINÁRIO
Senior Cyber Security Engineer
INOV - INESC Inovação Instituto de Novas Tecnologias (INOV)

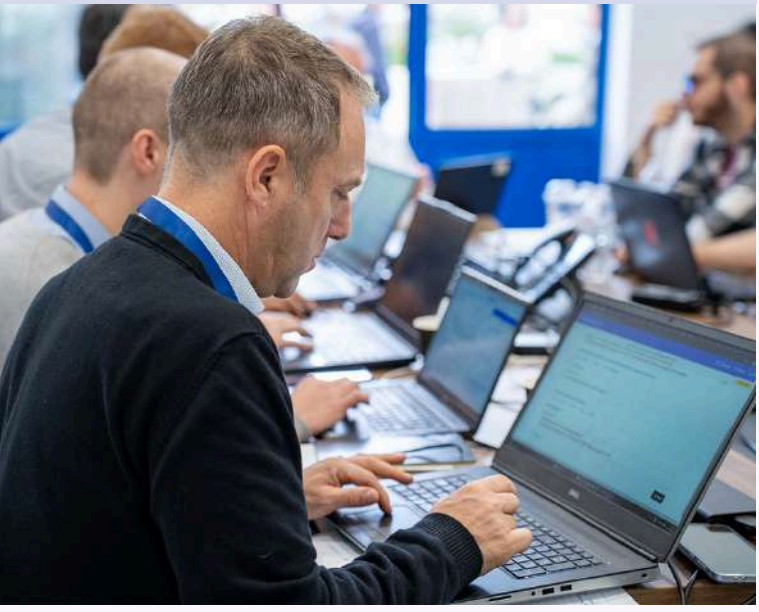

The STARLIGHT's co-development approach has been central to turning research concepts into solutions that genuinely support frontline investigative work. Rather than developing technology in isolation, the project embedded continuous dialogue with operational practitioners, allowing their needs, constraints, and real investigative scenarios to guide every stage of development. This ensured that the tools evolved in response to actual challenges encountered in the field, not theoretical assumptions.

The two tools presented in the following pages, Logo Detection and the Cyber Pattern Investigator (CPI), reflect this process in practice. Each represents a collaboration that unfolded over multiple cycles of feedback, testing, refinement, and validation with dedicated law enforcement partners.

Through this sustained engagement, the tools matured significantly, achieving levels of usability, performance, and relevance that align with day-to-day investigative realities. Both tools address clearly defined operational gaps: one in video triage and visual identification, the other in cyber investigations where traditional digital evidence may be unavailable. Their development illustrates how co-creation with practitioners not only strengthens the technical quality of the tools but also ensures that the outcomes can be directly applied within investigative workflows.

Together, these stories demonstrate the value of STARLIGHT's model: practical innovation grounded in operational insight, delivering tools that are ready to support law enforcement in tackling emerging and complex security challenges.

Logo detection, a tiny detail can become a big lead In many investigations, a seemingly insignificant logo on a jacket, cap, or rucksack becomes the detail that breaks open an identification. Yet until now, detecting such logos across hours of video required investigators to manually scan footage frame by frame, a slow, eye-straining process ripe for improvement. STARLIGHT's Logo Detection Tool, developed by the Centre for Research and Technology Hellas (CERTH) in collaboration with TNO, changes that dynamic. It automatically detects brand logos and emblems in both live video streams and large archives of investigative footage, turning what was once a labour-intensive task into a fast, structured workflow.

What sets this tool apart is not only its accuracy, but its accessibility. Training an AI model for a new logo typically requires deep technical knowledge. Under STARLIGHT, CERTH and TNO engineered a system that hides this complexity behind a simple interface, allowing investigators to add new logos with only a few clicks. Achieving this balance, a powerful model training beneath an intuitive user experience, was one of the project's major technical challenges.

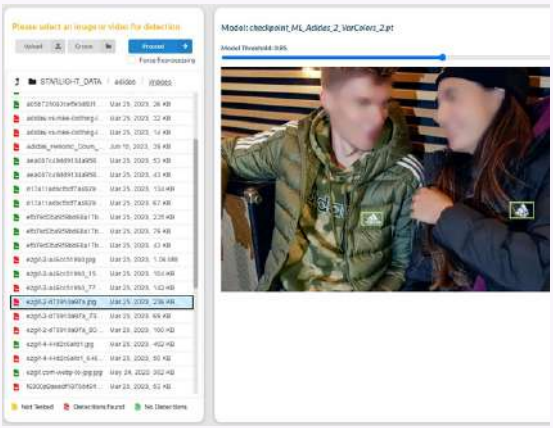

Close collaboration with law enforcement partners, particularly the Spanish Ministry of Interior (ESMIR), shaped the tool's evolution. LEA feedback led directly to features such as batch processing, improved training workflows, and support for live-stream analysis. These refinements ensured that the tool reflects investigative reality rather than purely technical assumptions. The tool is now advancing under the EACTDA Tools4LEAs programme, where it will continue to grow in maturity and operational alignment. For investigators dealing with ever-growing volumes of video data, the Logo Detection Tool offers a smarter, faster way to turn a tiny visual clue into a meaningful lead.

1. Enables investigators to train models for detecting new logos without AI expertise
2. Works on both archival video and live camera streams
3. Features shaped directly by LEA input
4. Addresses a recognised operational gap in video triage

The Cyber Pattern Investigator (CPI), finding ransomware traces when evidence is gone

Ransomware attacks rarely leave investigators with the luxury of intact digital evidence. Disks may be encrypted, wiped, or unavailable, forcing analysts to rely on the fragments that survive: network traces buried deep inside firewall logs. Extracting useful intelligence from these logs traditionally required specialist knowledge and many hours of manual review.

The CPI tool, co-developed by INOV - INESC Inovação Instituto de Novas Tecnologias (INOV) and Polícia Judiciária (PJ), bridges this gap. It provides law enforcement with a structured, semi-automated way to identify Indicators of Compromise (IoCs) hidden in firewall logs, even when disk evidence is gone. CPI enriches raw logs with embedded cyber intelligence such as geolocation, protocol context, cloud and VPN indicators, and then applies AI-driven clustering and statistical outlier detection to highlight suspicious patterns. With just a log upload, investigators receive a clear, visual, explainable report designed specifically for scientific police teams. The result: investigative hours reduced to minutes, and a clearer path to understanding how an intrusion may have unfolded.

The tool's development was driven by continuous co-creation between INOV and PJ's UNC3T cybercrime unit, whose operational insights shaped everything from interface design to workflow structure and the forensic reporting format. Regular ToolFests and pilots helped validate CPI on both synthetic and real case data, ensuring that the tool reflects practical investigative needs. Delivered as a Dockerised dashboard for easy deployment, CPI is expected to continue evolving beyond the end of STARLIGHT through ongoing collaboration and training. In a time where cybercriminals increasingly destroy digital traces, CPI gives investigators a valuable new ability: turning firewall logs into a meaningful forensic trail.

1. Extracts IoCs from firewall logs when disk evidence is unavailable
2. Uses AI clustering and outlier detection to reveal hidden patterns
3. Co-developed and validated with PJ cybercrime specialists
4. Deployable across agencies through a Dockerised dashboard

LEAs reflections: over four years of co-developing AI for European policing

Luisa PROENÇA
Deputy National Director
Polícia Judiciária (PJ)

Manuel ÁLVAREZ & Álvaro GALISTEO
Spanish Ministry of Interior (ESMIR)

Paul WOUTERS
Coordinator of Strategic Analysis
Belgian Federal Police (BFP)

Over the 52 months of STARLIGHT, law enforcement agencies across Europe worked side by side with researchers and developers to co-create a new generation of AI tools for policing. Rather than developing technology in isolation, the project adopted a co-development (CODEV) approach that placed investigators at the centre of every design choice.

Their insight, validation, and operational testing ensured that tools were not only innovative but realistic, trusted, and aligned with legal, ethical, and operational requirements. The result is a shared legacy: tools that address concrete investigative gaps, a strengthened understanding of AI governance, and a collaborative pathway for the responsible adoption of AI in European law enforcement.

The perspectives shared in this article reflect contributions from three Law Enforcement Agencies (LEAs) actively engaged in STARLIGHT's co-development cycles.

Their experiences illustrate the broader patterns observed across all fifteen LEAs involved in the project.

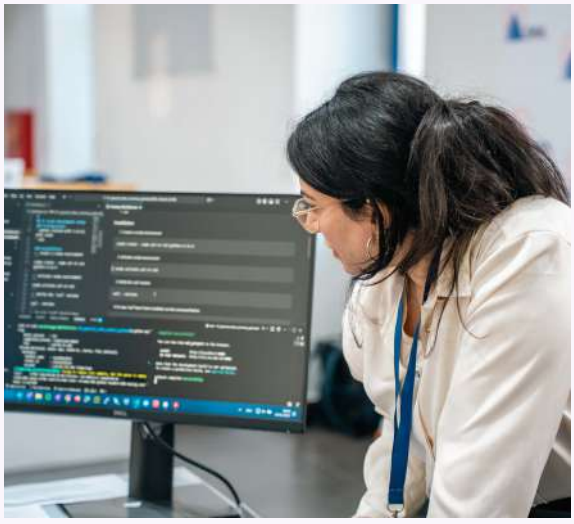

From expectations to real-world impact At the start of STARLIGHT, some law enforcement agencies had only limited exposure to AI systems, while others were already familiar with emerging capabilities. In every case, expectations evolved considerably as agencies gained hands-on experience.

The Belgian Federal Police (BFP) recalled having "a general view of AI's potential," but through STARLIGHT, "gained deeper insight into the broad range of AI possibilities and policing applications covering the shared needs of LEAs, as well as into its growing performance."

The Spanish Ministry of Interior (ESMIR) entered the project with previous experience in AI technologies yet explained that STARLIGHT broadened their appreciation of what was realistically achievable, leading to a perspective that deepened throughout the project.

Polícia Judiciária (PJ), Portugal, likewise noted that their understanding shifted as the tools matured, stating that STARLIGHT "allowed us to see how AI could directly support complex investigative tasks in a way we had not seen before."

Across all agencies, the progression was clear: AI moved from theoretical possibility to tangible operational support.

Tools that filled real investigative gaps One of the clearest messages from LEAs is that STARLIGHT delivered tools with capabilities not found in commercial platforms. PJ highlighted this most strongly in relation to the Cyber Pattern Investigator (CPI) prototype, describing it as "a tool built entirely from ideas and needs expressed by PJ experts," and noting that several of its capabilities "are not available in any commercial market solution".

ESMIR emphasised the practical value of STRACT, a tool that substantially reduced the time required to analyse data gathered through covert technical means. They also reported considerable benefit from the Logo Detection tool, which fills the gap of forensic searches across videos and images based on a single reference image, and from the Audio Denoiser, which serves as a valuable first step toward improving tactical audio recordings collected in complex scenarios.

BFP stressed the usefulness of tuning AI technologies to provide practical support for law enforcement tasks today. Benefits are concrete and immediate, tools that reduce workloads and improve efficiency rather than theoretical demonstrations of technical capability.

Real examples of added value The agencies shared specific examples demonstrating where STARLIGHT tools made a measurable difference. PJ validated CPI using real closed cases, showing how the tool could surface relevant connections that would otherwise require extensive manual work and confirming that the results produced were the same as those achieved by manual analysis.

ESMIR used STRACT to analyse geolocation data, identifying patterns, matching target positions, and even forecasting their potential future movements. They also confirmed that the Logo Detection tool proved highly effective when training models for targeted searches across extensive video archives.
These examples illustrate not only how the tools function, but how they enhance real investigative workflows.

Trust, explainability, and human oversight All law enforcement partners expressed a shared view that AI must support, not replace, investigators.

ESMIR captured this most clearly: "human oversight is essential: AI cannot replace the judgement of investigators". BFP emphasised that AI tools must provide an appropriate level of transparency and control and enable users to understand, question, compare, and challenge their outcomes. PJ noted that STARLIGHT's human-in-the-loop approach corresponds closely with investigative reality: "in investigations, decisions must remain human, even when AI accelerates the analysis".

Partners widely recognised the value of STARLIGHT's ethical foundations, including the AP4AI principles, the project's focus on embedding legality and proportionality, and the deliberate separation between machine-generated suggestions and human decision-making. These elements were instrumental in building confidence in adopting AI within sensitive investigative settings.

CODEV cycles: a new model for AI development
Across all feedback, the CODEV process emerged as one of STARLIGHT's most important contributions. Instead of presenting fully formed tools at the end of development, STARLIGHT involved LEAs from the earliest stages, working together through iterative cycles of design, testing, and refinement.

PJ emphasised that CPI was built entirely from ideas and needs from PJ experts, describing the process as uniquely aligned with operational practice. ESMIR highlighted that their feedback directly influenced tool features, describing this close collaboration as rare in research projects and deeply valuable. The co-development approach stands as one of STARLIGHT's lasting scientific and operational contributions.

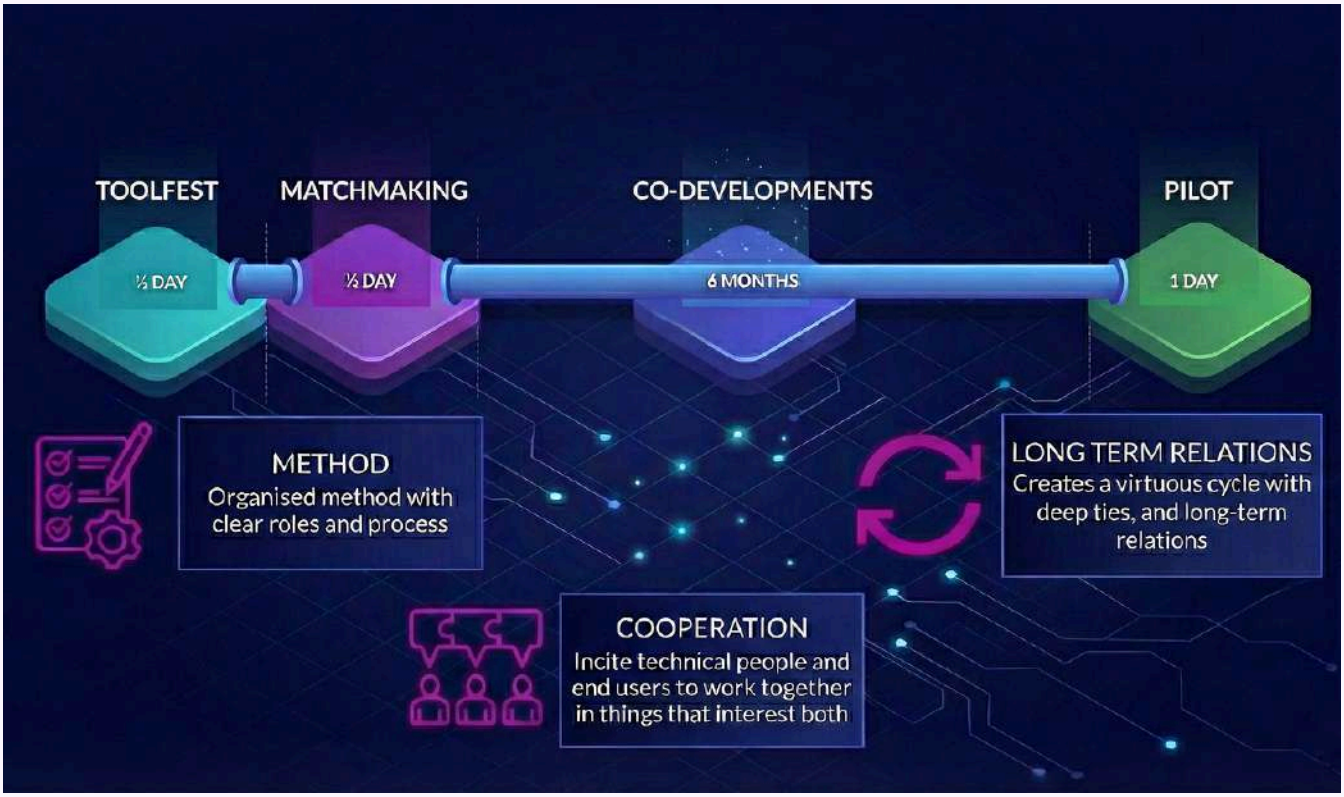

The STARLIGHT Co-development Methodology

What LEAs will carry forward Looking beyond the project, LEAs consistently underscored that STARLIGHT's most enduring legacy lies not only in its tools but in the skills, methods, and governance approaches developed throughout the four years of collaboration. These capabilities remain within agencies irrespective of the lifespan of individual prototypes.

Agencies highlighted a strengthened methodological capacity: the ability to validate AI tools, understand limitations, and appropriately interpret outputs within casework. This competence is now embedded in internal investigative and analytical routines.

They also noted a higher level of ethical and governance readiness. STARLIGHT gave partners practical experience in applying responsible AI principles, defining appropriate levels of human oversight, and becoming aware of their responsibilities in deploying AI systems, all of which support alignment with emerging regulatory requirements such as the AI Act.

Perhaps the most transformative legacy is cultural. LEAs expressed a clear intention to continue co-developing technologies rather than adopting solutions built in isolation. The iterative, collaborative approach tested in STARLIGHT has become recognised as a framework for future development and procurement, an approach that keeps tools grounded in operational reality.

STARLIGHT's lasting contribution When asked about STARLIGHT's long-term impact, LEAs converged on a common view: the project did not merely deliver tools, it reshaped the way AI is developed and adopted in European policing.

BFP summarised it succinctly: "STARLIGHT's most significant contribution is pointing the way forward to building an ecosystem for aligning innovation with law enforcement adoption. This is essential for responding to the ultimate challenge, i.e. bringing innovative AI developments to a TRL9 and achieving operational use by LEA end-users." ESMIR reflected that the project strengthened their organisation's capacity to adopt AI in a controlled and ethical way, while PJ emphasised that STARLIGHT provided solutions that align with investigative realities and the legal context in which they operate.

STARLIGHT's legacy is therefore multi-layered: a portfolio of operationally validated AI prototypes, a strong ethical and governance foundation, a collaborative model for future development, increased readiness for the AI Act, and renewed trust between practitioners and developers.

Together, these contributions help position European law enforcement for a future in which AI is used responsibly, effectively, and always with human judgement at its core.

Sustaining the STARLIGHT legacy: building a lasting future for responsible AI

> Our aim is to make STARLIGHT's legacy part of a broader ecosystem, one where research, policy, and operational needs converge to strengthen Europe's resilience and security.
>
> **Juan ARRAIZA IRUJO, Executive Director, EACTDA**

As STARLIGHT reaches its conclusion, attention turns toward sustainability, ensuring that the tools, datasets, and communities developed throughout the project remain accessible and operational for the European security ecosystem.

The project was designed to enable Law Enforcement Agencies (LEAs) to integrate responsible AI tools into everyday work. Using gap analysis and hands-on experimentation through ToolFests, LEAs have helped define real operational needs and shaped tools that respond to them.

Turning results into operational capability A central element in achieving this continuity is the collaboration with the European Anti-Cybercrime Technology Development Association (EACTDA). Through its Tools4LEAs initiative and the Security Research Results Repository (SRRR), we will ensure STARLIGHT's results will continue to benefit LEAs across the EU and Schengen area.

Planned to launch in early 2026, the EACTDA Portal will serve as a single gateway for eligible public entities fighting cybercrime to access trusted, cost-free technologies, including those initially developed within STARLIGHT. Three tools are already being transitioned and are currently undergoing "last-mile" development to reach full operational readiness.

EACTDA's approach guarantees sustainability through open licensing, ongoing maintenance, and early-adopter support.

Complementary exploitation pathways While EACTDA provides a confirmed route for further development and distribution, STARLIGHT's sustainability strategy also includes complementary channels. Europol's Tool Repository (ETR) will provide access to mature STARLIGHT results under its established licensing framework, strengthening LEAs' operational capabilities across Europe.

STARLIGHT's ToolFests supported this vision by connecting developers and end-users in direct collaboration. They served as a launchpad for new cycles of co-development, ensuring that STARLIGHT's innovations continue to evolve in line with real-world needs.

> EACTDA plays a central role in ensuring this transition, supporting the long-term operationalisation of STARLIGHT results through the Tools4LEAs initiative and the SRRR.
>
> Juan ARRAIZA IRUJO, EACTDA

Sustaining STARLIGHT's communities and data
Beyond technical outputs, STARLIGHT has built vibrant communities of expertise in areas such as counter-terrorism, Child Sexual Abuse Material (CSAM) prevention, cybercrime, legal and ethical compliance, and DevSecOps. We have ensured collaborations to maintain these networks and their knowledge exchange.

While EACTDA will not directly maintain the STARLIGHT platform, both organisations are exploring ways to ensure interoperability through API-level integration, shared datasets, and potentially single sign-on access. STARLIGHT's datasets and AI models may also be hosted in EACTDA's Security Research Results Repository (SRRR), an open-access platform for sharing research outcomes and models.

A lasting ecosystem for responsible AI The collaboration between STARLIGHT and EACTDA demonstrates how European research can transition into long-term operational capability. Through shared governance, interoperable platforms, and strong community engagement, STARLIGHT's legacy will continue to grow beyond the project's lifetime, strengthening Europe's resilience, fostering responsible AI adoption, and bridging the gap between research, policy, and law enforcement practice.

Together, these mechanisms ensure that STARLIGHT's impact extends well beyond the project's lifetime. Through EACTDA's operational framework, Europol's established platform, and continued engagement with LEAs, STARLIGHT's legacy will endure.

By embedding responsible, explainable, and trustworthy AI tools into operational practice, the project has helped lay the foundation for a more resilient and adaptive European security landscape. The journey does not end here, it evolves into a broader ecosystem that unites research, policy, and practice in the pursuit of a lasting future for responsible AI in law enforcement.

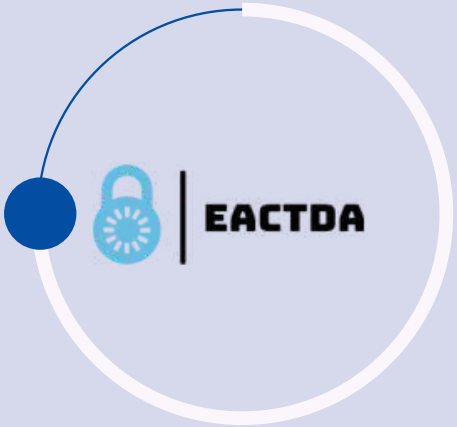

STARLIGHT

Sustainable Autonomy and Resilience for
LEAs using AI against High Priority Threats

www.starlight-h2020.eu
starlight@cea.fr