Automating Detection & Classification for Mine Countermeasures on AUVs

A Cathx White Paper

Overview

Mine countermeasure (MCM) operations remain among the most demanding and resource intensive missions in naval and security domains. Current practice typically involves a twostage process: an autonomous underwater vehicle (AUV) conducts a wide-area sonar survey to detect potential threats, followed by a separately launched uncrewed surface vessel (USV) and a remotely operated vehicle (ROV) to reacquire and visually classify the contacts. This approach requires multiple vehicle deployments, sustained operator involvement, and mission durations measured in days rather than hours.

Through this process, confirmed live threats are ultimately neutralized, while inert or low-risk objects are often left in place and recorded in naval hydrography databases as hazards to avoid. But this process takes time, and any delay increases risk. The central bottleneck lies not in detection, but in the transition from sonar detection to reliable visual classification.

Advances in autonomy, perception, and real-time decision-making technologies now present an opportunity to collapse the timeframe of the operation. By enabling AUVs not only to detect, but also to visually classify objects of interest in real time, operators can accelerate timelines, reduce risk, and increase mission efficiency under the demanding conditions of littoral environments.

Cathx is developing the **Clarity MCM software module** to address this challenge.

Clarity provides an intelligent perception layer for AUVs sensors, fusing multiple data sources with real-time processing and model execution onboard the vehicle to enable real-time identification and classification of objects. Instead of simply collecting raw data for postmission review, Clarity extracts critical information from raw data in real time, delivering meaningful inputs to the vehicle upon which it can respond dynamically to objects of interest – whether mines, infrastructure, or anomalies. For manufacturers, Clarity provides a way to accelerate autonomy development by rapidly integrating advanced perception and classification capabilities without the need to build them from scratch.

Autonomous MCM is a Force Multiplier

The following compares MCM mission times for surveying a 20.8 sq. km area and acquiring ground-truth visual imagery for 20 sonar-identified contacts.

Detailed mission time breakdowns are provided in Appendix 1.

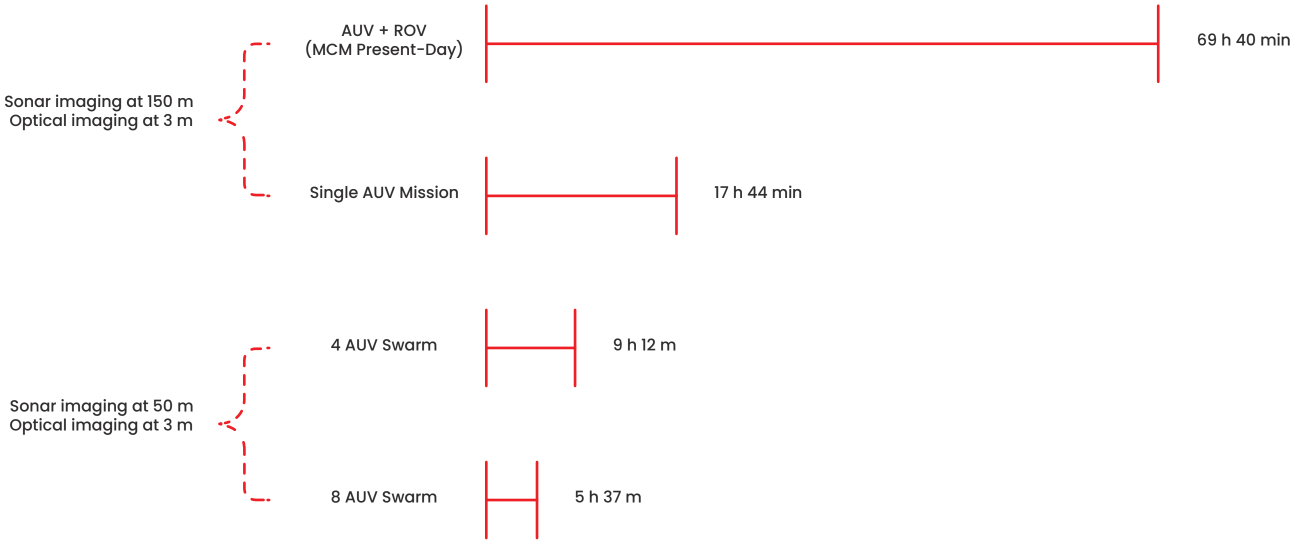

The Core Challenge

Uncrewed systems have changed what's possible in the field. But for MCM, real-time classification remains the limiting factor. Workflows are still slower and more resource-heavy than they should be.

Present-day missions typically unfold as follows:

- An AUV is deployed to conduct a wide-area sonar survey, often over 20+ km² of seabed at operational depths of several hundred meters.
- On recovery, the sonar data is processed and analyzed ashore, leading to the identification of potential contacts of interest.
- Each of these contacts must then be reacquired and classified to determine whether it represents a mine or a non-mine object.
- To achieve this, a small ROV is launched from a surface vessel, usually with two crew onboard (though some operators now use uncrewed platforms), and is guided back to each location.
- The ROV conducts reacquisition passes using multibeam sonar before descending into camera range (~3-7 meters above the seabed) to capture the optical imagery required for positive visual identification.

This process is slow , operationally complex, and places personnel and platforms at sustained risk.

The core challenge, and the focus of this paper, is how to move from today's multi-step model to a single autonomous AUV mission that can both detect and visually classify contacts in real time. Our approach achieves that goal while also improving current workflows. By solving key perception and processing challenges on the vehicle, it lays the groundwork for future autonomy developments, including dynamic mission planning.

Scope Note: This paper addresses the role of autonomous systems in the detection, classification, and identification phases of mine countermeasure (MCM) operations. It focuses on how emerging autonomy, perception, and real-time decision-making technologies can accelerate and de-risk these critical steps. The methods and technologies described here are not concerned with the subsequent disposal or neutralization of mines, which remains outside the scope of this document.

The Role of Cameras in Autonomous MCM

As AUVs take on roles traditionally carried out by ROVs and human operators, the importance of a high-quality onboard camera increases significantly. In conventional MCM workflows, an ROV pilot visually inspects sonar contacts and gathers the visual evidence needed for final classification and decision-making. In an autonomous workflow, this responsibility shifts to the AUV itself.

A high-quality camera is therefore essential to ensure that every contact is captured photographically. These images provide the ground truth required for final verification by a human operator, maintaining confidence in the system's findings.

Equally important is the ability to process images in real time. An onboard image classifier ensures that sonar contacts are captured by the camera as the vehicle passes over them. Without this step, there is a risk of incomplete coverage, which could force costly and time-consuming second missions. By confirming contact capture during acquisition, classifiers help guarantee mission completeness and make the overall workflow significantly more efficient.

In this way, the camera becomes more than a supporting sensor; it is a critical enabler for autonomy, bridging the gap between machine-driven detection and the human oversight still required in operational MCM.

Dynamic Mission Planning & the Perception Gap

Most AUVs today still operate along pre-set survey lines, collecting data and avoiding obstacles with limited adaptive behavior. To move beyond this, vehicle manufacturers are investing heavily in next-generation autonomy engines designed to support dynamic mission planning and, in some cases, more advanced functions such as swarm control. These engines represent a step change in capability: instead of executing static plans, vehicles can adjust their behavior in the field, re-plan tracks in response to new information, investigate potential contacts, and optimize area coverage in real time. This transition is essential for enabling AUVs to shift from being data collection platforms to becoming active decision-making systems that can carry out complex missions with reduced operator intervention.

However, the effectiveness of these autonomy engines ultimately depends on a capability that is still underdeveloped: real-time onboard perception. For dynamic mission planning to work as intended, the vehicle must be able to process and interpret sonar and optical data as it is collected and then feed those insights directly into the autonomy engine. Today, most AUVs still push raw data back to operators for post-mission analysis, which means the autonomy layer lacks the critical inputs it needs to make timely decisions mid-mission. Closing this gap requires solving challenges in onboard processing, bandwidth, and multi-sensor data fusion. Until those issues are addressed, autonomy engines will remain limited in practice, unable to deliver the full promise of responsive autonomy.

Enabling Detection & Classification on the AUV

Executing both detection and classification within a single AUV mission represents a fundamental advance in MCM capability. Eliminating the need to recover survey data for shore-based processing, redeploy additional vehicles, and revisit targets streamlines the workflow into a single, continuous operation.

This approach reduces vessel and crew requirements, compresses mission timelines, and accelerates decision-making in contested environments. While an operator-in-the-loop remains essential for final decision authority, single-mission execution ensures rapid delivery of high-quality imagery to support timely and confident threat classification.

Technical Requirements for On-Vehicle Detection and Classification

Achieving the goal of conducting both detection and classification within a single mission depends on meeting several technical requirements. These include the development and deployment of ATR and classification models, the handling of environmental and operational constraints, and the limits of onboard systems. Specifically, the technical requirements include the following:

- Fine-tuning detection performance for ATR of sonar data
- Deploying sonar ATR to the AUV and establishing an effective data workflow
- Collecting sufficient imagery to train object classification models
- Deploying visual classifier models to the AUV and establishing an effective data workflow
- Operating in conditions of poor visibility (littoral zones and tidal environments)
- Managing power consumption within mission constraints
- Bandwidth limitations for communicating images required for human-in-the-loop verification

Different Vehicle Patterns Emerging

AUV manufacturers are exploring different design patterns for delivering autonomous MCM capability. One possibility is a single-AUV mission, where both sonar and camera payloads are integrated into the same vehicle. The main appeal of this approach is simplicity: with all sensors and processing in one platform, operations could be easier to plan, deploy, and manage without the added complexity of coordinating multiple vehicles.

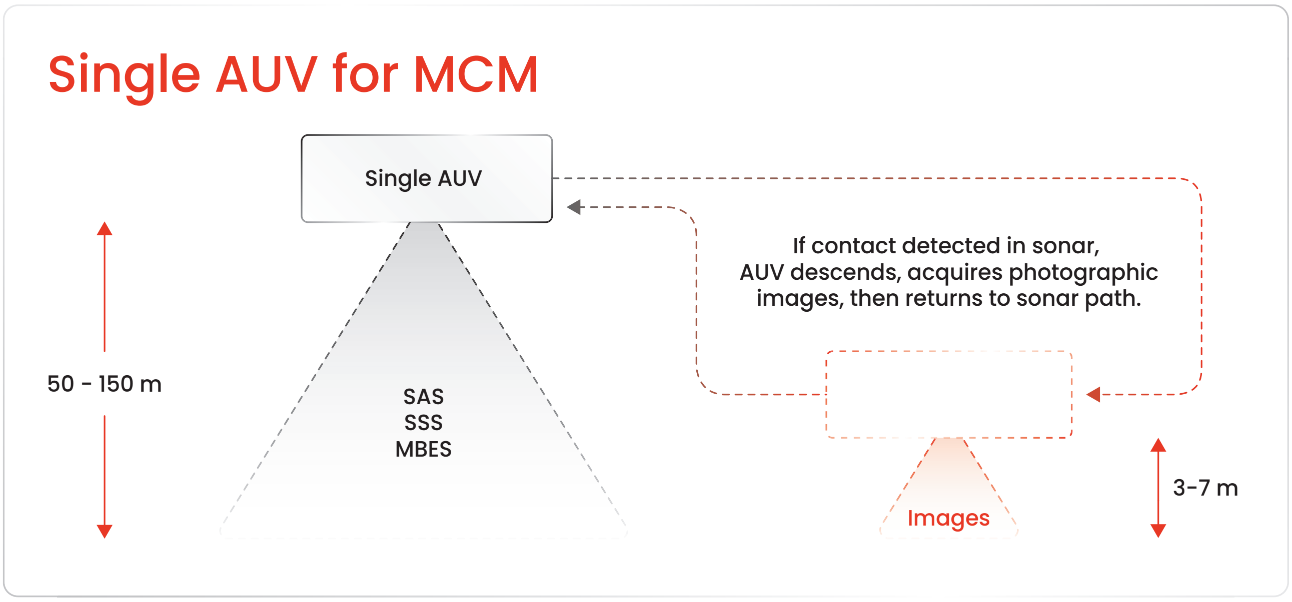

Another possibility is a **paired-AUV mission**, in which one vehicle is equipped with sonar and communicates detections to a second vehicle carrying a camera. The advantage here lies in specialization: each vehicle could be optimized for its role, which may enable more efficient sensing and even parallel operations—for instance, sonar mapping continuing while a second vehicle investigates contacts.

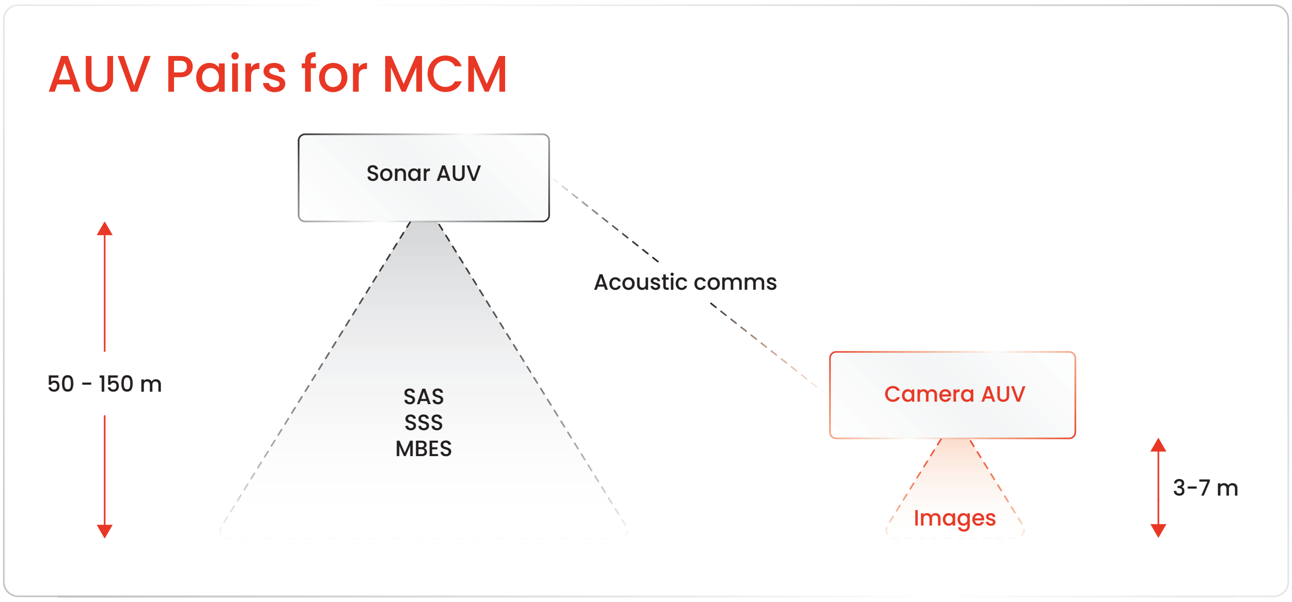

Both approaches present interesting trade-offs, and it is not yet clear which will become dominant. The direction taken will likely depend on evolving mission requirements, operational constraints, and the pace of development in autonomy engines capable of coordinating sensor data and vehicle behavior in real time.

Implications for Swarm Operations

Looking further ahead, each pattern may also influence how swarm operations develop. A single-AUV approach offers simplicity but may scale less efficiently if every vehicle must carry the full suite of sensors. A paired-AUV approach could allow roles to be distributed across a swarm - for example, some vehicles focused on detection, others on classification . This division of labor might make multi-vehicle missions more efficient, but it also introduces new challenges in coordination and control. These are early ideas, and it remains an open question which model, or combination of models, will prove most effective for autonomous MCM.

Automating Detection & Classification on AUVs

Cathx has released the following core capabilities to the market:

- Edge processing hardware (INKA) for AI-based data processing on any vehicle
- Edge workflow software (Clarity) to turn sensor data into actionable information
- Optical sensors for high-performance imaging

The Cathx INKA Iris is a compact imaging system designed for smaller AUVs with integrated INKA hardware for executing Clarity data processing workflows, including AI model inferencing. INKA Iris makes it easy to deploy AI models and build data processing workflows for execution on AUVs. Events detected in Clarity workflows can be communicated to the AUV autonomy engine to support dynamic mission planning. The AUV autonomy engine can trigger Clarity workflows using the Clarity API.

AUV Solution Architecture for Sonar Detection and Image Classification on the Edge

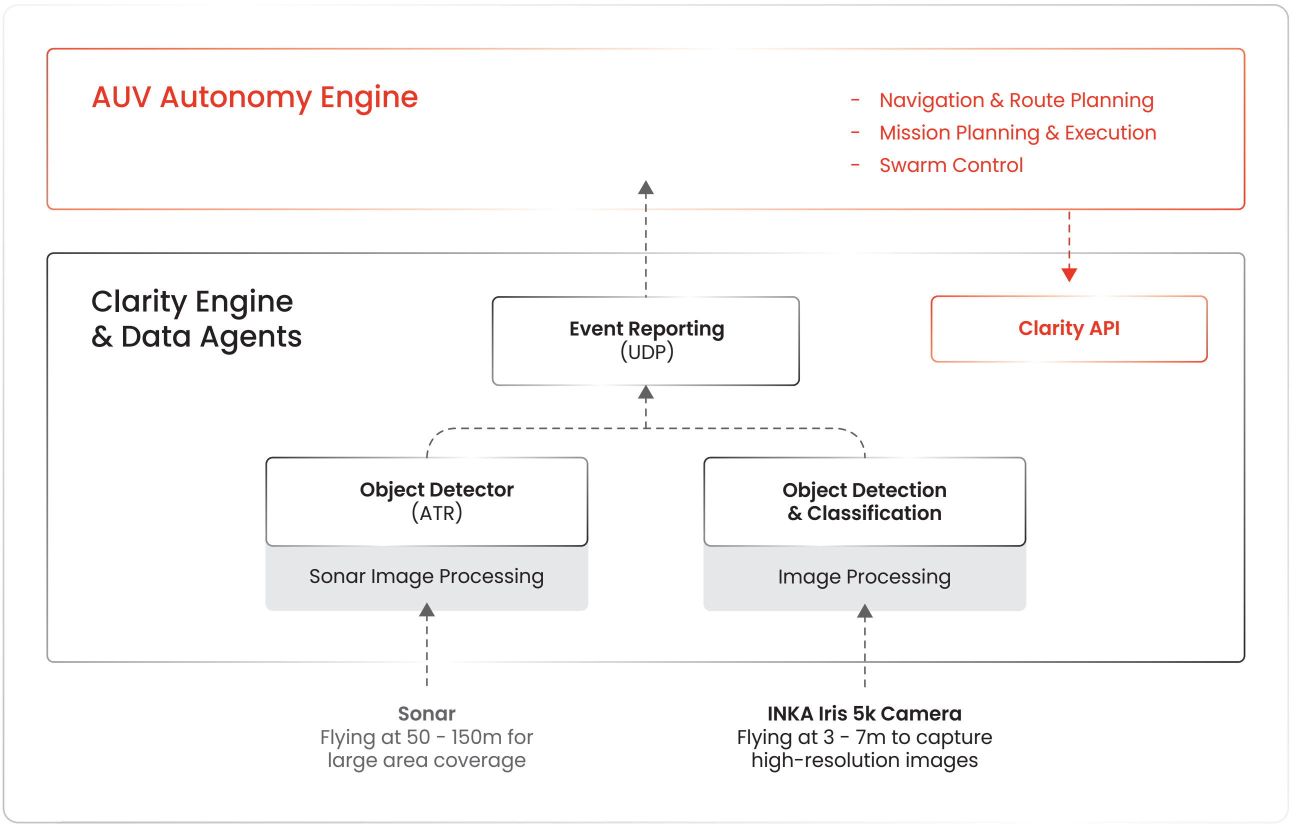

Cathx is working with our customers to leverage the INKA Iris and Clarity Data Agents (modular data processing blocks) to meet the technical requirements for developing autonomous MCM detection and classification capabilities.

CATHX

Autonomous MCM Requirements	Cathx Solution / Approach
1. **Fine-tuning model performance for automatic target recognition (ATR) in sonar data**	Building on our experience in visual object detection, Cathx is developing a general object detection model for sonar data as a Clarity Data Agent. This serves as a baseline ATR capability that customers can further refine and evolve with their own mission-specific training data. 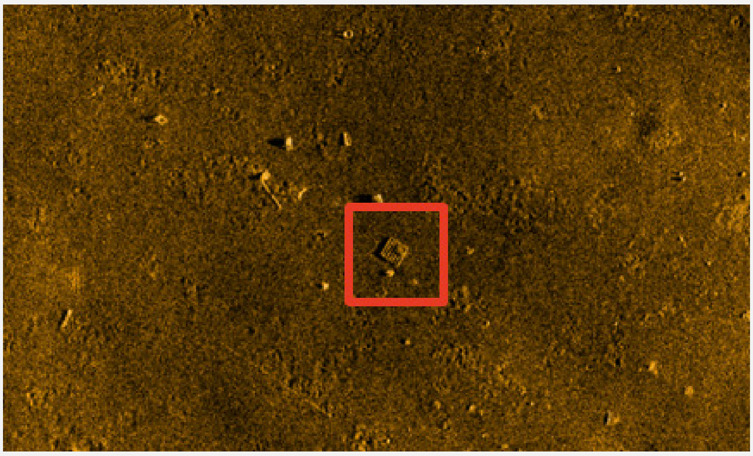
2. **Deploying sonar ATR to AUV and creating data workflow**	Integrating INKA Iris onto an AUV delivers the ability to run Clarity data processing workflows on data as it is being acquired. Each Clarity workflow is a sequenced combination of one or more Clarity Data Agents, blocks of code that execute a data specific task. Clarity Interface Data Agents can be built and configured to work with 3rd party sensors. Clarity workflows are being developed to take processed sonar images, run additional image processing steps (as required), then pass the sonar image to the Cathx ATR model. An alert or event message can then be sent by Clarity to the vehicle autonomy engine.

Autonomous MCM Requirements	Cathx Solution / Approach
3. Collecting images to train object classification model	INKA Iris enables the capture of large volumes of high-resolution images for creating labelled training data and improving model performance. INKA Iris comes with a set of Clarity Data Agents for image processing, a critical step to producing high-quality, standardized images and improving model performance.
4. Deploying visual classifier model to AUV and creating data workflow	The Clarity AI model inferencing Data Agent enables AI models to be deployed and run as part of a Clarity workflow. It is also possible to deploy retrained and updated models back to the vehicle on subsequent missions without the need to reinstall or update any Clarity software.
5. Poor visibility Littoral zone, tidal conditions	INKA Iris synchronized strobe lighting (400,000+ lumens) dramatically reduces motion blur and reduces the impact of turbidity. Combined with the on-sensor Clarity Data Agents for image processing, INKA Iris delivers ultra-sharp images, even in harsh shallow-water environments.

Autonomous MCM Requirements	Cathx Solution / Approach
6. **Power consumption**	INKA Iris operating at 3 FPS and running a full Clarity AI-based workflow would require ~45 W, as illustrated in the following table: 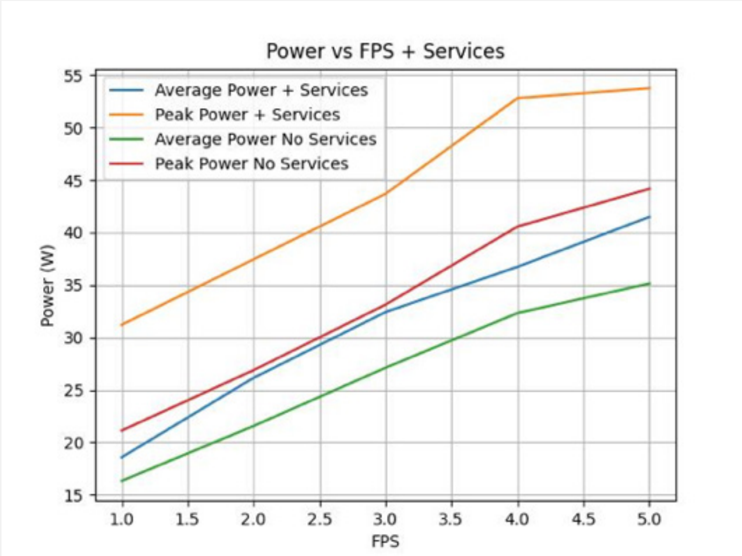
7. **Bandwidth** Communication of images for human-in-the-loop verification	Dynamic workflows can be executed using the Clarity API to select the optimal image format for the available AUV communication bandwidth. Further, by running the image classification model as part of a Clarity workflow, only relevant images are transmitted.

Note on Customer Models and IP: All models, datasets, and intellectual property developed by customers remain exclusively owned by the customer. Cathx has no visibility of these components at any stage. Clarity is designed to operate in fully isolated environments, ensuring that customer workflows, models, and mission data are deployed and executed locally on the vehicle without requiring Cathx access or oversight.

Getting Started

The end goal for autonomous mine countermeasures (MCM) is the capability to complete the full cycle of detection and classification in sonar and image data in a single mission, without reliance on post-mission data processing and analysis. Achieving this requires incremental development pathways that address current capability gaps in autonomy, automatic target recognition (ATR), and end-to-end data workflows.

Cathx provides both the hardware (INKA Iris with integrated edge processing hardware) and the software (Clarity workflows with modular Data Agents) to enable these capabilities as progressive steps. By working in partnership, we can accelerate the transition from today's multi-vehicle, multi-day missions to single-mission autonomy.

The Cathx approach does not remove the human operator from the final stage of visual confirmation and human judgment remains essential in verifying contacts of interest. INKA Iris and Clarity provide an intelligent perception layer for AUV sensors, fusing multiple data inputs with real-time processing and model execution onboard the vehicle to enable real-time identification and classification of objects. This significantly reduces the volume of raw data to be analyzed and shortens the timeline from initial survey to operator confirmation from days to hours.

At Cathx, we recognize that each manufacturer and operator is at a different stage in their journey toward true autonomy. Our technology is designed to be modular and scalable, enabling vehicle manufacturers and operators to adopt capabilities in a way that best aligns with their mission requirements and development timelines. The following pathways allow customers to adopt Cathx technology at the pace and scale that best fits their operational needs.

Pathway 1: Better Images for Model Training

There is typically a shortage of high-quality optical imagery, data that is critical for training and improving the performance of machine learning models for object classification. Building larger, higher-quality image datasets is therefore a key first step toward training robust object classifiers.

INKA Iris directly addresses this challenge. Designed for easy integration on AUV platforms, INKA Iris makes it straightforward to begin acquiring the optical data needed to strengthen training pipelines .

Pathway 2: Developing an ATR Model

Where no ATR capability exists, Cathx recommends a prototype development approach. Using the Cathx ATR baseline model as a foundation and combining it with additional training data provided by the vehicle manufacturer, we can generate a working ATR prototype. Importantly, the training data does not need to include real mines; replica objects placed on the seabed are sufficient for model development. This approach allows rapid iteration, validation, and deployment of an ATR capability as a Clarity workflow. The result is a functional ATR model that can evolve in fidelity with additional datasets.

Pathway 3: Adding Perception to the AUV

A critical requirement for autonomous operation is the ability to run perception workflows directly on the vehicle. INKA Iris addresses this with integrated edge hardware and software capable of executing AI models in situ. Within the Clarity framework, detection algorithms and machine learning models are deployed as Data Agents, allowing perception tasks to operate natively on the AUV without adding external processing overhead.

Equally important, Clarity workflows can be structured to provide real-time inputs to the autonomy engine. By linking perception outputs, such as detections, classifications, or confidence scores directly into mission control logic, the vehicle can adapt its behavior dynamically. Clarity's API framework enables autonomy engines to call custom workflows directly, while workflows themselves can send detections or event messages back into the autonomy engine. This creates a continuous feedback loop in which perception and planning operate as a single system. Such tight coupling between sensing and control is central to closing the loop from data acquisition to autonomous decision-making and forms the foundation for responsive MCM operations.

Pathway 4: Custom Integrations and Workflows

Working in partnership with Cathx, customers can:

- Integrate third-party sensors such as sonar (MBES, SAS, SSS, and SBP), magnetometers, or custom navigation systems for use in Clarity workflows.
- Tailor workflows to mission objectives.
- Continuously evolve workflows, adding new detection algorithms, retrained AI models, or mission-specific processing steps over time.

This makes it possible to implement advanced autonomy features such as adaptive survey planning, on-the-fly sensor fusion, and selective data capture based on mission context.

By combining modularity with extensibility, Cathx allows customers to build mission-optimized autonomy, without the overhead of developing an entirely new software architecture from scratch.

Appendix 1: Auv Scenarios & Time Estimations

Present-Day MCM Mine Hunting Scenario & Time Estimate

In this mine countermeasures scenario, an AUV is deployed to conduct an eight-hour sonar survey, mapping 20.8 square kilometers of seabed at an average operating depth of 300 meters. On completion of the mission, the vehicle is recovered, and its sonar data is processed and analyzed ashore, leading to the identification of approximately 20 potential contacts of interest.

Each of these contacts must be reacquired and classified to establish whether it represents a mine or a non-mine object. To achieve this, a small ROV is launched from a two-person surface vessel and guided back to each location. The ROV first conducts reacquisition passes at 50 meters altitude using multibeam sonar, before descending into camera range to capture the optical imagery required for positive visual identification.

The time estimates are indicative and can vary based on environmental conditions, seabed complexity, navigation and reacquisition accuracy, and processing workflows.

Table 1: Present-Day MCM Identification Process and Time Breakdown

Mission 1: AUV using sonar to identify contacts and areas of interest		
Step 1	**AUV Launch** Plan flight path and launch AUV.	2 h
Step 2	**AUV mine hunting survey using sonar** The AUV conducts an eight-hour survey at 150 m altitude and 3 knots, covering 20.8 km² and producing ~2.5 TB of raw sonar data.	8 h
Step 3	**Sonar Data Offload and Quality Check** Transfer the collected sonar data from the AUV and conduct QA/QC analysis to ensure data integrity before processing.	2 h

Step 4	**Sonar Data Processing and Mosaic Creation** Process the collected sonar data and generate linear mosaics to enable detailed target detection and classification.	24 h
Step 5	**Target Detection and Classification** Automated target recognition (ATR), supported by human-in-the-loop analysis, identifies 20 contacts requiring optical verification.	8 h
	Mission 1: AUV Total time	**44 h**

Mission 2: ROV to confirm contacts using optical data		
Step 6	**ROV Deployment to Contact Site** Transit the vessel to the first contact location, then prepare, launch, and navigate the ROV to 300 m depth for reacquisition.	45 min
Step 7	**Contact Reacquisition** Conduct one to two passes at 60 m altitude using multibeam sonar with a 120 m swath to reacquire the designated contact.	10 min
Step 8	**Visual Classification of Contact** Approach the contact with the ROV and capture high-resolution photographic imagery for visual confirmation.	5 min
Step 9	**ROV Transit to Next Contact** Navigate the ROV from the current site to the next designated contact location to continue the verification sequence.	15 min
	Time for optical verification for 1 contact	1 h 15 min
	Mission 2: ROV Total Time (20 contacts)	25 h
	Total mission time (AUV + ROV)	**69 h**

Timeline for AUV and ROV MCM Operation:

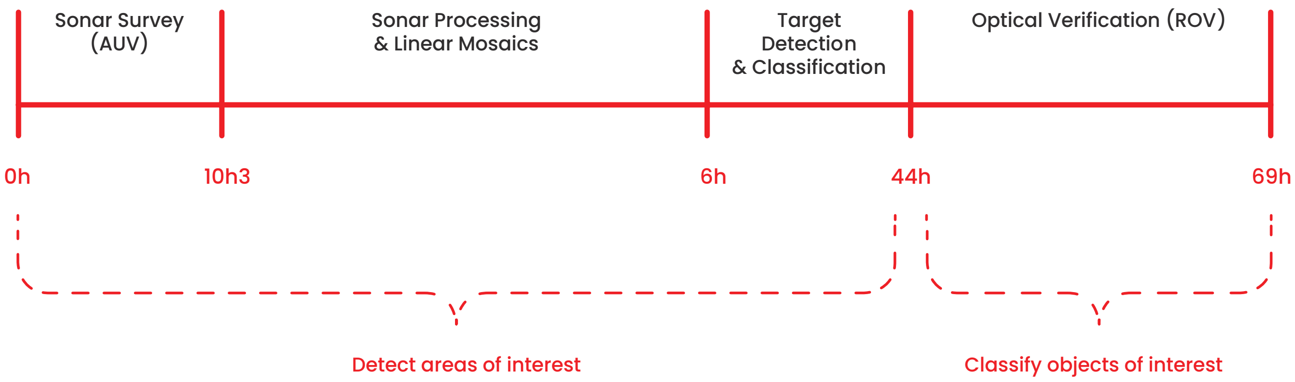

Single-AUV Mission Scenario & Time Estimate

In this MCM hunting scenario, a single AUV is tasked with completing both detection and classification functions in a single mission. The vehicle conducts an eight-hour sonar survey at an operating altitude of 150 meters, transiting at 3 knots to map an area of 20.8 square kilometers. Onboard ATR using sonar data is used to identify potential contacts of interest in near real time.

For each identified contact, the AUV descends from 150 meters to the seabed to capture optical imagery. Imaging is conducted over a 10 m × 10 m area with 40% side overlap, ensuring sufficient ground-truth coverage necessary for accurate visual capture and classification of the contact. After imagery is collected, the vehicle ascends back to 150 meters and resumes its survey track.

The AUV is equipped with a Cathx INKA Iris with embedded ATR capabilities which, in addition to capturing high-resolution imagery at 5 FPS, incorporates onboard edge-processing hardware capable of executing data workflows to run both the ATR detection model and the image classifier model. Notably, these models do not run simultaneously, and because all processing is performed within the INKA Iris, there is no additional computational load or data-processing demand placed on the AUV itself.

Table 2: Single AUV Mission Process and Time Breakdown

Single AUV mission		
Step 1	**AUV Launch** Plan flight path and launch AUV.	2 h
Step 2	**AUV mine hunting survey using sonar** The AUV conducts an eight-hour survey at 150 m altitude and 3 knots, covering 20.8 km². On-vehicle Sonar ATR running on INKA Iris in near real time.	8 h
Step 3	**Visual Classification of Contact** AUV approaches the contact at 3 knots and captures photographic imagery for visual confirmation. Visual Classification of Contact – Time Budget • Descent (150 m → seabed @ 15° pitch): ~5 min • Imaging a 10 m × 10 m area with 40% side overlap at 3 m altitude would require 7 lanes and 6 turns: ~1 min 20 sec • Ascent (seabed → 150 m): ~5 min Total time per contact: ~11 min 20 sec Total time for 20 contacts: 3 h 50 min At 3 FPS, images acquired per contact: 114 Images acquired for 20 contacts: 2280	3 h 50 min

Step 4	**Surface and Transmit Images** Assume 10 of the 20 contacts identified by sonar are classified on the AUV as mines. For each positive classification the AUV will ascend to the surface (assume 300 m depth) and transmit 3 JPEGs to enable human operator to confirm threat classification. The AUV then descends back to 150 m to continue survey. Data volumes (with 20% overhead) 3 Thumbnails (~0.5 MB each): 1.8 MB total 3 JPEGs (~3 MB each): 10.8 MB total Ascent time to surface at 15° pitch: 13 min Transmission times over Certus 350 • 3 Thumbnails: ~45 seconds • 3 JPEGs: ~4 minutes Time to descend back to 150m: 6.4 min Total time per contact image transmission cycle (assume 3 JPEGs are sent): 23.4 mins	3 h 54 min for 10 contacts visually classified as mines
	Total Mission Time (Single-AUV)	**17 h 44 min**

Autonomous MCM Using AUV Swarms

Note: AUV swarm capabilities remain at a relatively early stage of development, and the scenario described here represents just one of many approaches currently being explored by different vehicle manufacturers and operators.

In this MCM scenario, an XL-class AUV acts as a “mother” platform, deploying and recovering a swarm of small AUVs. Each small vehicle is equipped with a multibeam echosounder (MBES) and a Cathx INKA Iris optical imaging system. Operating at an altitude of 50 meters and a speed of 3 knots, each AUV achieves a usable sonar swath of approximately 170 meters with 40% overlap, enabling the swarm to rapidly cover large areas in parallel. We assume 20 contacts evenly distributed over an area of 20.8 square kilometers. Each AUV ascends to 3 meters to acquire images and visually classify each contact when it is identified by the sonar ATR.

Table 3: AUV Swarm Scenario and Time Breakdown

		4 AUV Swarm	8 AUV Swarm
Step 1	**AUV Launch** Launch small AUVs from XL AUV.	20 min	40 min
Step 2	**AUV mine hunting survey using sonar** Each AUV executes a survey track allocated by the XL AUV, while onboard sonar ATR, processed in near real time on the INKA Iris, enables local detection of contacts.	8 h 10 min	4 h 6 min
Step 3	**Visual Classification of Contact** • Descend 50 m → 3 m: ≈2 min • Image 5 m × 5 m @ 3 m, 40% overlap • Coverage: 3 lanes × 6 frames/lane = 18 frames ≈22 sec • Ascend back to 50 m: ≈2 min Total per contact: ~4.4 minutes	22 min	11 min
Step 4	**AUV Recovery** Recover AUVs to XL AUV	20 min	40 min
	Total AUV Swarm Mission Times	**9 h 12 min**	**5 h 37 min**

sales@cathxocean.com
Ireland: +353 (0)45 252786
United States: +1 (281) 374 1476
United Kingdom: +44 (0)1224 432 180
Australia: +61 867 278 886

www.cathxocean.com