First Mile Transparency & Traceability on the Edge

INTEGRATING TRACEABILITY AND AI-POWERED MONITORING TO DELIVER ACTIONABLE INSIGHTS ON VERIFICATION, RISK, AND QUALITY

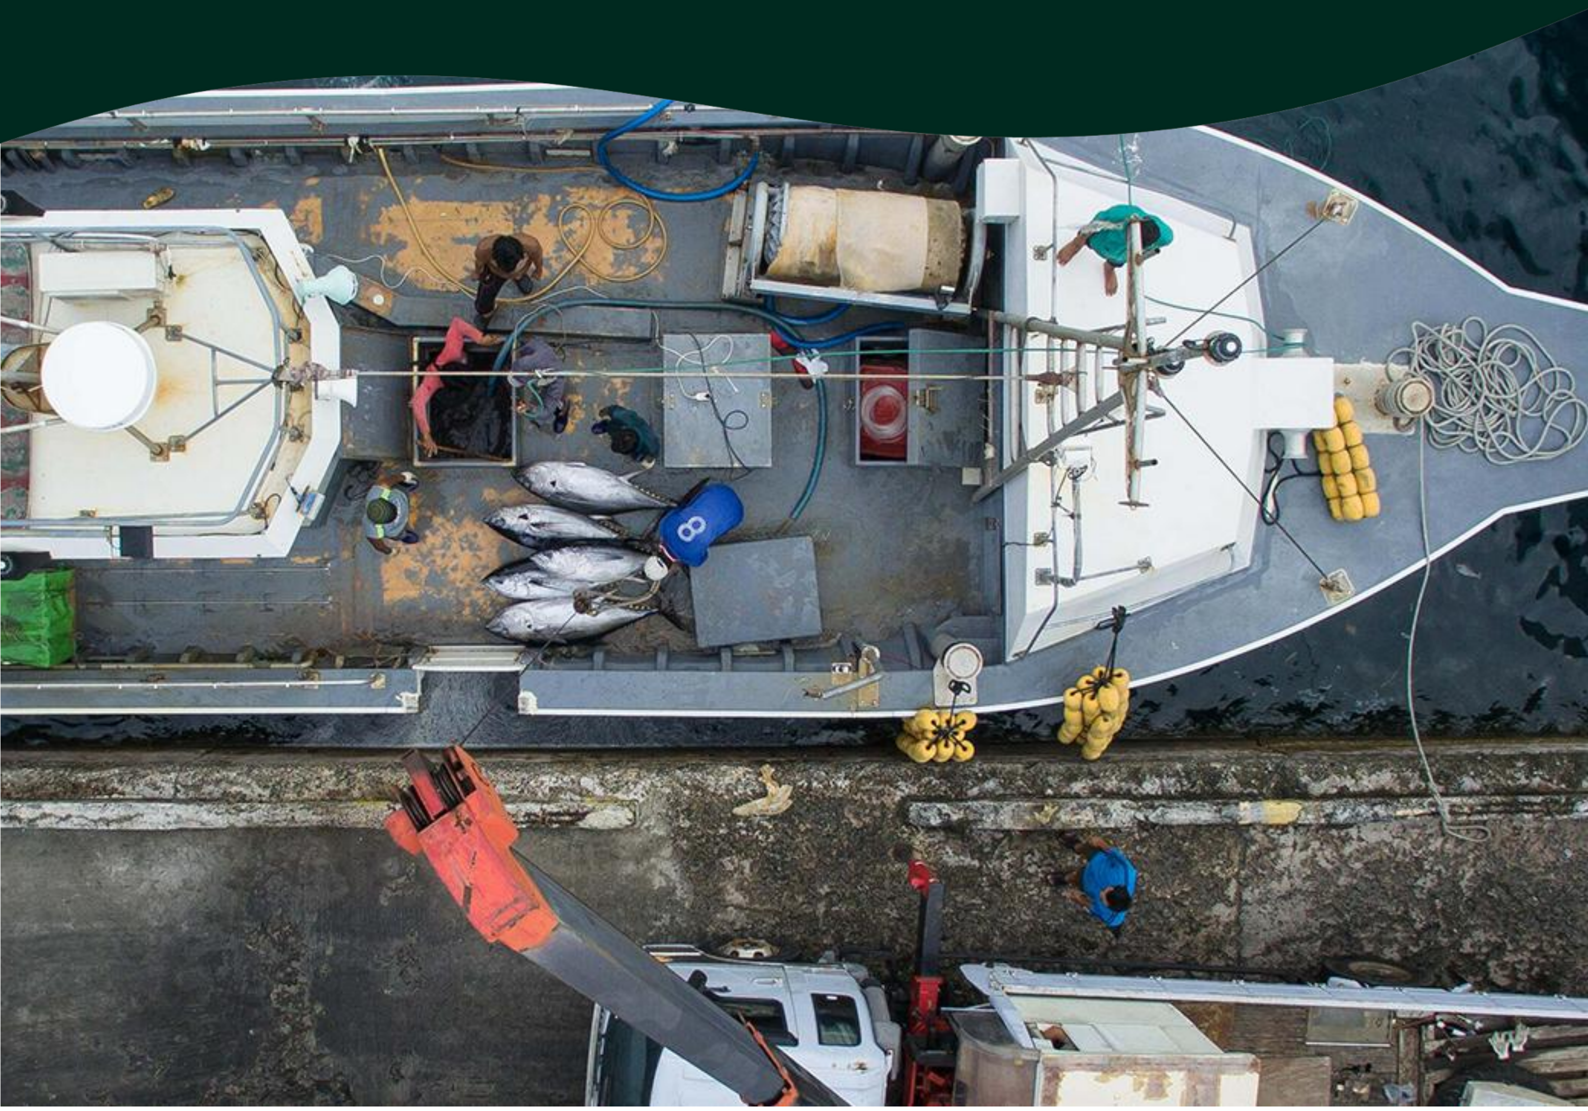

PREPARED BY

tryo·labs

PREPARED FOR

The Nature Conservancy

Acknowledgements

The Tryolabs team would like to thank the network of partners whose collaboration made this project, led by The Nature Conservancy, possible.

We sincerely appreciate the commitment and support of our project partners: InMotion, Wholechain, Bureau Veritas, Deckhand, Thalos, the Sea Pact members, the Instituto Costarricense de Pesca y Acuicultura (INCOPESCA), and the Cámara Nacional de Industria de Palangre (CNIP).

We are especially grateful to the vessel owners, captains, and crew who generously shared their time, expertise, and effort in the field. Their participation was essential to the success of this work.

This work was made possible thanks to The Nature Conservancy's collaboration with the Patrick J. McGovern Foundation and other supporters.

About Tryolabs

Tryolabs is an AI and data partner that helps organizations turn data and AI ambition into production-ready systems. With 15+ years of experience, we work with companies and institutions across industries to design and deploy custom machine learning solutions that are robust, scalable, and built to perform in the real world. Our expertise spans computer vision, MLOps, and edge AI, with a track record of shipping systems that operate reliably in complex, resource-constrained environments

Suggested citation

The Nature Conservancy and Tryolabs are sharing this work to help support the wider adoption of first-mile traceability solutions that link quality and sustainability information to individual fish—advancing efforts towards transparency and responsibility across seafood supply chains. Please cite this work using the following reference:

Tryolabs, The Nature Conservancy. (2026). First-mile traceability on the edge: *Integrating RFID tags, sensor data, and AI-powered monitoring for fish-level transparency* (Report prepared for The Nature Conservancy)

PHOTOS, FRONT AND BACK COVER: © Jonne Roriz

Foreword

Safeguarding the health of our oceans—and the communities and industries that depend on them—requires a new generation of tools that bring clarity, trust, and accountability to the seafood supply chain. That responsibility begins at the *first mile*, at the point of capture. For decades, processors, buyers, importers, and regulators have relied largely on unverified, self-reported catch documentation to determine market access. As a result, fisheries managers have lacked the scientific and operational data needed not only to manage fisheries sustainably, but also to ensure compliance with fisheries laws. Both fisheries managers and market actors have been forced to make critical commercial and enforcement decisions without timely, verifiable information about what happens at sea or under what conditions fish move from capture to market. Fundamental questions often go unanswered: Was this product harvested legally, and can it be purchased with confidence? Should a vessel be inspected based on the activities that occurred during a given trip? What did the cold chain look like, how long can the product be expected to remain on the shelf, and was it comingled with catch from other vessels?

This initiative marks a meaningful shift in what is possible.

By merging AI-powered electronic monitoring (EM) with fish-level traceability, we demonstrate a technology stack capable of transforming conservation outcomes, commercial insight, and industry performance. Building on our work to operationalize AI-enabled EM systems onboard longline vessels, this project extends that foundation into the first mile of the supply chain by linking each fish to where it was caught, how it was handled and stored, and the relative risk associated with each fishing trip.

For conservation, the implications are profound. Near real-time species-level catch data and trip risk scores give managers the intelligence needed to move from reactive oversight to proactive stewardship. For industry, the benefits are equally compelling. By capturing handling conditions and linking them to individual target catch, the system provides practical indicators of product quality—signals that producers can use to differentiate their catch and protect against product comingling. Complex operational data are distilled into intuitive quality and sustainability indicators, and cumbersome paper-based catch documentation is digitalized at the vessel to enable efficient access downstream in the supply chain. This system strengthens trust between producers and buyers and creates the enabling environment for markets to benefit from transparency assurances and reward responsible practices.

Importantly, this work shows that advanced electronic monitoring and traceability can be integrated into industrial fishing operations without adding significant burden on fishing crews. The system's modular, interoperable design ensures compatibility with existing on-deck workflows and traceability platforms, lowering barriers to adoption and enabling broader, more equitable access. Furthermore, our solution is open access for the larger seafood technology community. To learn more about the AI-powered electronic monitoring system, see the report titled "*Monitoring Fishing Activity at the Edge.*"

As environmental pressures intensify and demand for sustainable seafood grows, the need for transparent, verifiable, and actionable information has never been greater for companies and regulators. The technology demonstrated here—AI-powered EM paired with fish-level traceability—offers a practical, scalable path toward ensuring a world where healthy oceans and thriving seafood economies advance together.

Ben Gilmer
Director
Large-Scale Fisheries Program

Executive summary

This project piloted a first-mile seafood traceability system on a commercial longline vessel. During fishing operations, crew members tagged individual fish, allowing each catch to be digitally linked to where it was caught, how it was handled onboard, and sustainability information generated by an AI-powered electronic monitoring (EM) system. This created a complete, fish-level traceability record that can be shared directly with seafood buyers.

The pilot deployment demonstrated several important outcomes:

First, sustainability scores derived from an AI-powered EM system (Tryolabs, 2026) can be mobilized for market and supply chain use, not just compliance. These scores were generated onboard, linked to specific fishing sets, and shared downstream as part of the traceability record, enabling clearer product differentiation and more transparent buyer–supplier relationships.

Second, meaningful product quality proxies can be derived reliably from onboard sensor data and associated with tagged fish. Time-to-ice, median ice-well temperature, distance to coast, and trip duration were captured automatically and combined to characterize handling and storage conditions during the first mile of the supply chain.

Third, complex traceability data can be translated into clear, user-friendly signals. Quality and sustainability performance were summarized into two intuitive traffic-light scores at the set level, allowing buyers to assess performance at a glance while preserving access to underlying metrics. These signals also provide a foundation for feedback and training that can support improved handling practices and responsible fishing behavior onboard.

Fourth, the traceability system proved interoperable with existing buyer-facing traceability platforms. Traceability data were consolidated onboard and transmitted to Wholechain, an end-to-end supply chain traceability platform, confirming compatibility with systems buyers already use.

Finally, the pilot showed that first-mile traceability can be integrated into industrial fishing workflows without imposing a significant burden on crews. Tagging and data capture fit naturally into routine operations, and crew feedback indicated that the process was easy to adopt—an essential condition for scalability.

Overall, the project demonstrated a practical and scalable pathway for delivering verifiable first-mile traceability, strengthening transparency at the start of the seafood supply chain, and enabling better-informed purchasing decisions, greater accountability and stronger trust between producers and buyers.

Table of contents

1. Introduction

Industrial fishing plays a vital role in global food systems, providing a primary source of protein for billions of people and supporting millions of livelihoods worldwide. Yet, the sector faces persistent challenges in ensuring both sustainability and transparency. Over one-third of global fish stocks are currently fished at unsustainable levels, and illegal, unreported, and unregulated (IUU) fishing continues to threaten marine biodiversity and responsible fishing practices (FAO, 2020). For seafood buyers, these risks translate into uncertainty in the quality and sustainability of the fish they source, creating growing demand to verify catch sustainability before products enter global supply chains.

> Much of this uncertainty originates in the first mile of the supply chain—from harvest to portside offload—where limited data and visibility allow IUU activities and poor practices to go undetected. To address this gap, The Nature Conservancy (TNC), in partnership with technology collaborators, has been equipping fishing vessels with electronic monitoring (EM) systems and a complementary AI-powered EM system that analyzes EM footage in near-real-time. Using edge computing and computer vision, this system replaces the historically slow, manual review of EM footage with timely, automated reporting that provides species-level catch data and sustainability scores directly from the vessel, representing a significant advance in fisheries transparency and compliance capacity (Tryolabs, 2026).

Building on this foundation, this project introduces first-mile traceability at the individual-fish level. Throughout this report, the term traceability system refers to the integrated onboard system that combines an AI-powered EM system with additional traceability components, including radio frequency identification (RFID) tagging and ice-well temperature sensing. Together, these components link each market catch from the moment it is caught through storage onboard and ultimately to portside offload, where the product enters global seafood supply chains. The traceability system enables seafood buyers to access both quality proxies—such as trip duration, time-to-ice, and median storage temperature—alongside AI-generated sustainability scores associated with the fishing set. Together, these indicators provide a clear, actionable view of both handling practices and sustainability performance at the start of the supply chain.

The traceability system was piloted on a semi-industrial longline fishing vessel operating in the Eastern Tropical Pacific that partnered with TNC. This report documents the design and implementation of the pilot, including onboard hardware, data capture and integration, and delivery of traceability records to Wholechain, a blockchain-based technology company that provides end-to-end supply chain traceability.

2. Existing AI-powered EM system

Tryolabs developed and deployed an AI-powered EM system to analyze fishing activity directly onboard vessels. The system enables near-real-time monitoring of fishing operations and generates sustainability scores that can be linked to individual fish records via RFID tags. This sustainability information is integrated into the traceability system described in this report (see Tryolabs 2026 for a full technical description of the AI-powered EM system).

2.1 AI-powered EM system

The AI-powered EM system runs on a compact edge device installed on the vessel and processes EM footage as fishing occurs. It performs three core tasks: 1) detecting fish in the video, 2) classifying each detected fish by species, and 3) tracking its movement across consecutive frames. As shown in Figure 1, the AI-powered EM system uses a virtual line drawn across the vessel's deck to determine whether a fish is retained (brought onboard) or discarded (returned to the water after being brought onboard). Together, these functions enable the AI-powered EM system to keep a continuous record of every detected, retained, and discarded fish throughout the fishing day—along with its species ID, the time it was captured, and the vessel's location at that moment.

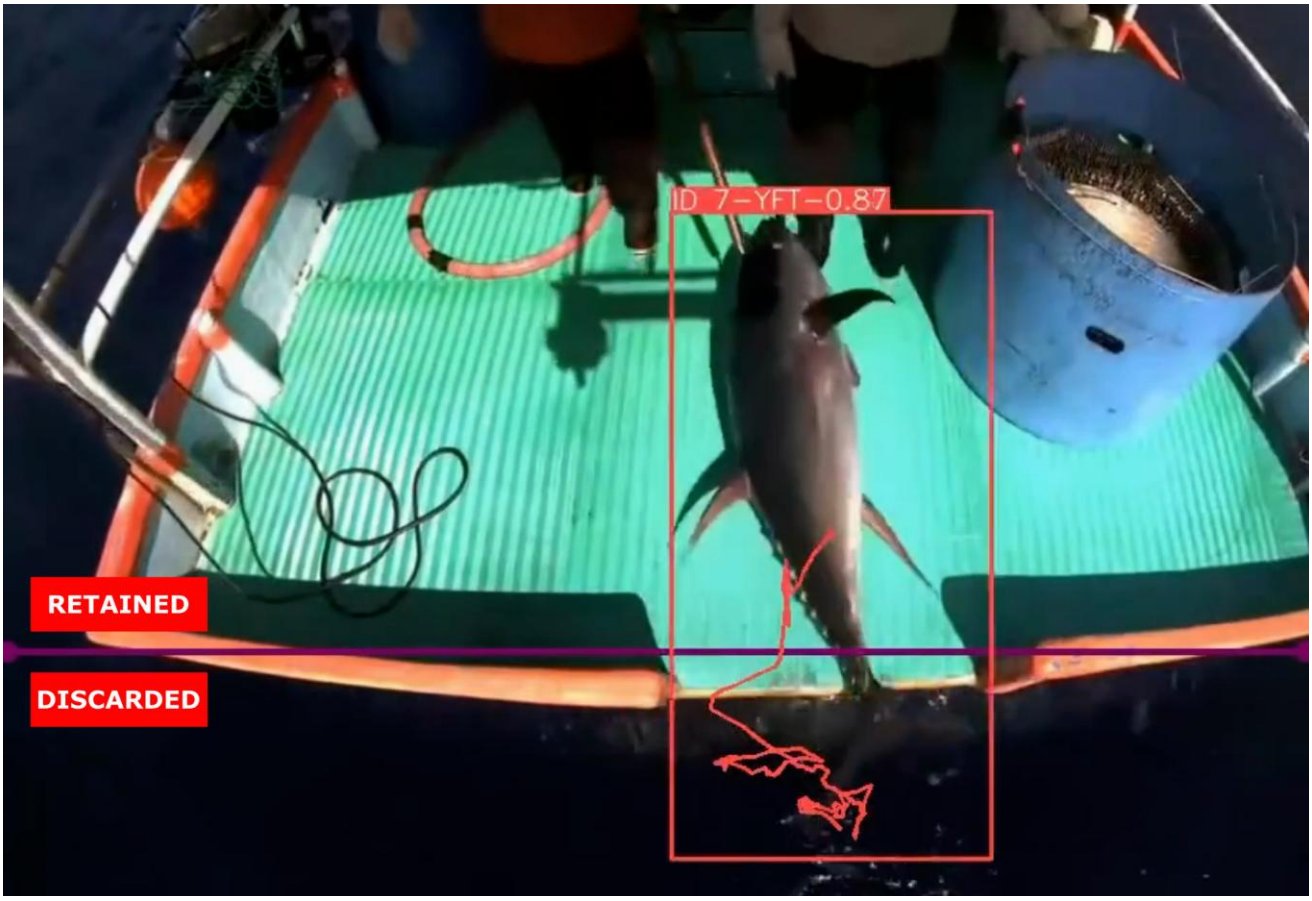

Figure 1. Example output from the AI-powered EM system showing fish detection, species classification, and tracking. A detected fish is shown in a red bounding box, labeled with a unique track ID, FAO species code (e.g., YFT for yellowfin tuna), and model confidence score. The red trajectory line shows the fish's movement across consecutive video frames. The purple virtual line represents a boundary on the deck used to classify events: tracks that cross onto the deck and remain onboard are labeled *Retained*, while tracks that cross and then return to the water are labeled *Discarded*.

2.2 Daily reports and sustainability scores

At the end of each fishing day, the AI-powered EM system compiles its observations into a structured *Daily Report*. As shown in Figure 2, reports include 1) a summary of retained and discarded catch counts; 2) sustainability scores that flag potentially non-compliant activity; and 3) species-level events with timestamps and EM video evidence clips.

A central element of the *Daily Report* is its sustainability score framework, which helps fisheries monitoring, control, and surveillance personnel prioritize secondary human review of fishing activity in EM footage. The AI-powered EM system assesses five daily risk variables, each highlighting a different type of potential concern during fishing activity (Table 1).

Table 1. Risk score components used in Daily Reports

Risk Component	Description	Primary Concern
Electronic Logbook Risk	Discrepancy between AI-predicted catch counts and captain's eLog records	Potential underreporting of catch by vessels
Illegal/ Sensitive Species Risk	Detection of sensitive species regulations prohibits	Retention of sensitive or illegal bycatch
GPS Risk	Geolocation of fishing activity	Fishing occurring in marine protected areas and other sensitive habitats
Model Underprediction Risk	Signs the AI may have missed catch events	Reduced reliability of AI-based monitoring
Operational Risk	Technical system issues such as missing GPS or video data	Loss of monitoring coverage that could enable non-compliance

Three of these components—Electronic Logbook Risk, Illegal/Sensitive Species Risk, and GPS Risk—are directly related to sustainability. Each is scored on a scale from 1 (low risk) to 3 (high risk). The traceability system aggregates these three values using a simple weighted average and maps the result to a traffic-light scale associated with each individually tagged fish to support intuitive interpretation (see Section 4.4 for more details).

2.3 Relevance for first-mile traceability

All fish caught on a given day inherit the same sustainability score, allowing each RFID-tagged fish to be linked to the sustainability conditions of the fishing activity it came from. Later in the workflow, this automatically-generated sustainability score is paired with a separate quality score—derived from onboard handling and ice-well storage indicators, as described in Section 4.4. Together, these scores give seafood buyers an at-a-glance view of both sustainability conditions and product handling during the first mile of the supply chain.

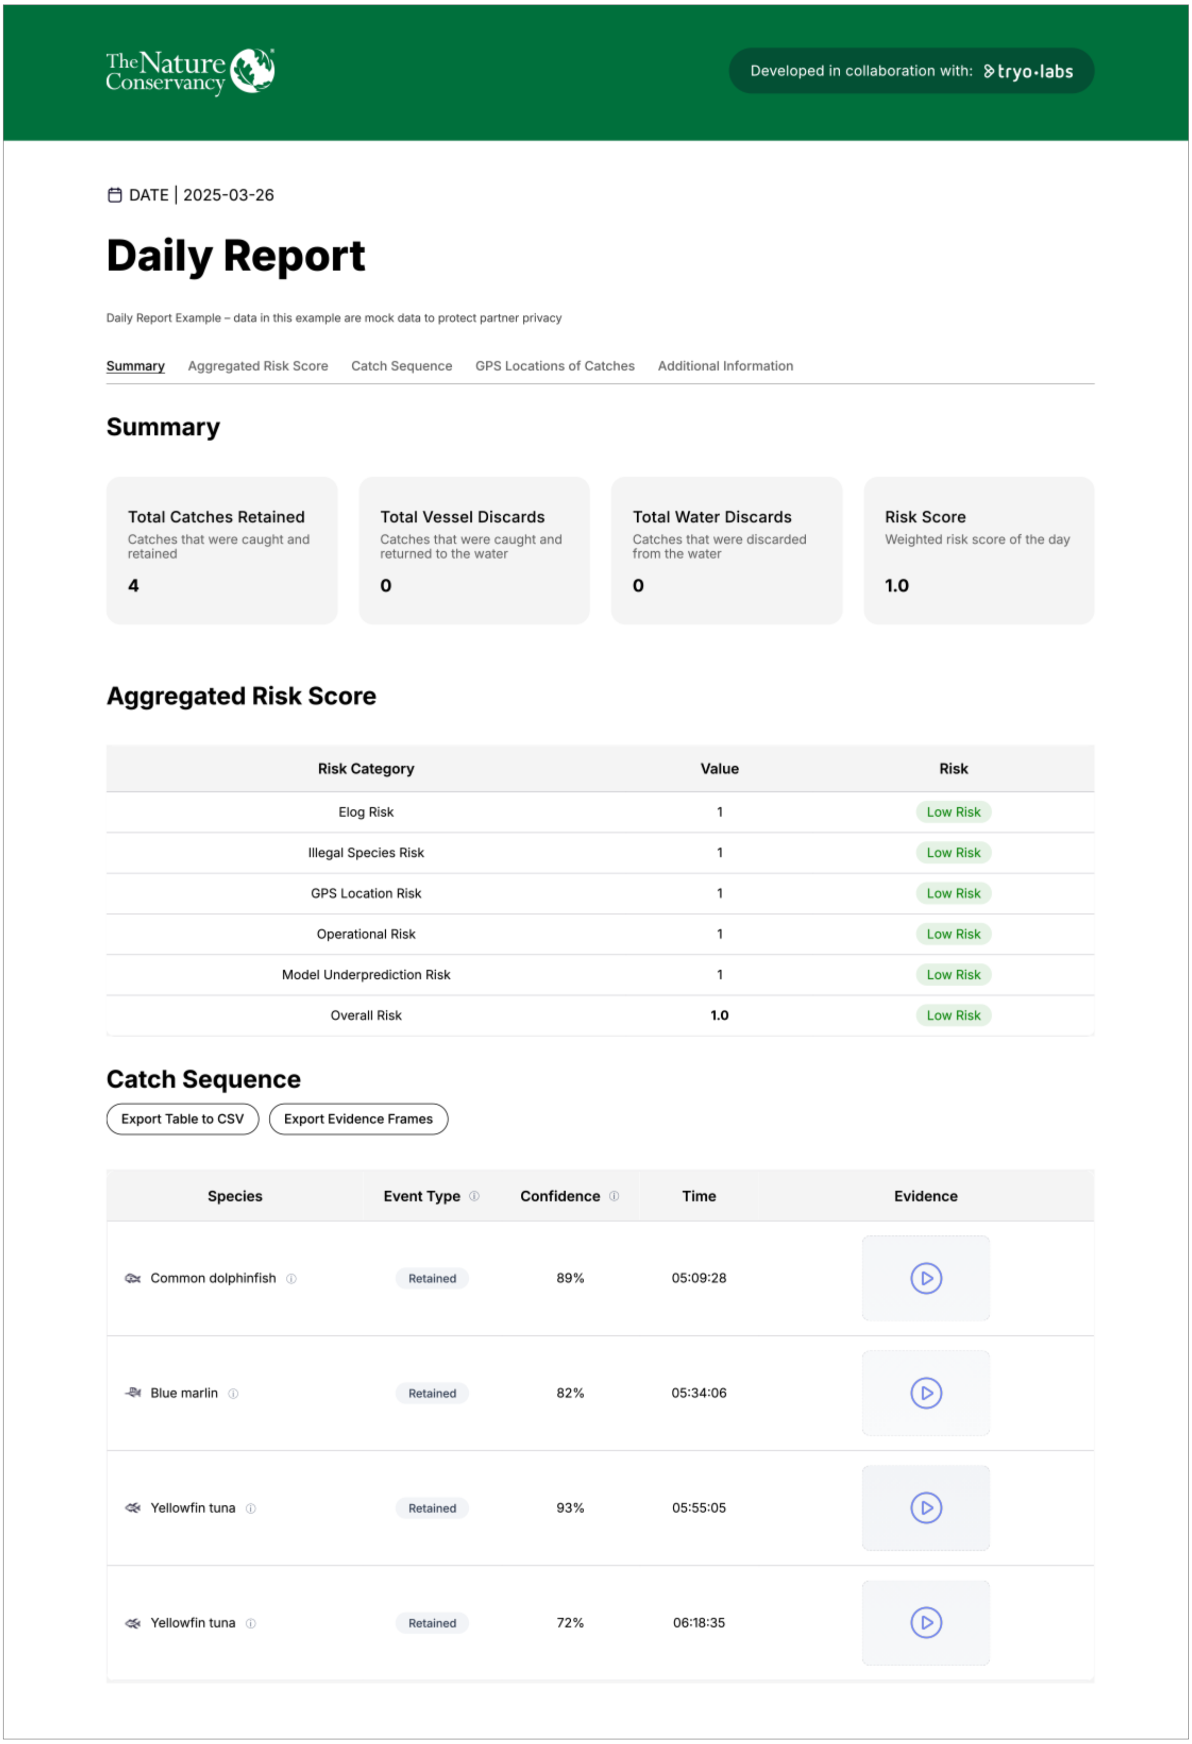

The Nature Conservancy

Developed in collaboration with: tryo·labs

DATE | 2025-03-26

Daily Report

Daily Report Example – data in this example are mock data to protect partner privacy

Summary Aggregated Risk Score Catch Sequence GPS Locations of Catches Additional Information

Summary

Total Catches Retained
Catches that were caught and retained
4

Total Vessel Discards
Catches that were caught and returned to the water
0

Total Water Discards
Catches that were discarded from the water
0

Risk Score
Weighted risk score of the day
1.0

Aggregated Risk Score

Risk Category	Value	Risk
Elog Risk	1	Low Risk
Illegal Species Risk	1	Low Risk
GPS Location Risk	1	Low Risk
Operational Risk	1	Low Risk
Model Underprediction Risk	1	Low Risk
Overall Risk	**1.0**	Low Risk

Catch Sequence

Export Table to CSV Export Evidence Frames

Species	Event Type ⓘ	Confidence ⓘ	Time	Evidence
Common dolphinfish ⓘ	Retained	89%	05:09:28	▷
Blue marlin ⓘ	Retained	82%	05:34:06	▷
Yellowfin tuna ⓘ	Retained	93%	05:55:05	▷
Yellowfin tuna ⓘ	Retained	72%	06:18:35	▷

Figure 2. Example of a *Daily Report* generated by the AI-powered EM system using mock data to protect partner privacy. The report includes: (1) a summary of retained and discarded catch counts; (2) a risk scoring table that flags potentially non-compliant activity; and (3) a detailed sequence of predicted catch events, each accompanied by species label, event type, timestamp, model confidence, and an EM video clipcaptured at the moment of counting.

3. Traceability workflow and hardware setup

The traceability system is designed to record fish-level handling, quality, and sustainability conditions during the "first mile" of the seafood supply chain—from the moment a fish is brought onboard until the vessel returns to port. This stage of fishing is typically the least visible segment of the supply chain, yet it is where many of the most important quality- and sustainability-related events occur.

The system focuses on capturing a small number of high-value events automatically, minimizing additional workload for the crew while providing a reliable foundation for downstream traceability. The workflow below shows how these events are captured during routine fishing operations.

3.1 Traceability workflow

Figure 3 provides an overview of the traceability workflow, showing how routine fishing operations are paired with traceability data capture and AI-generated sustainability scores.

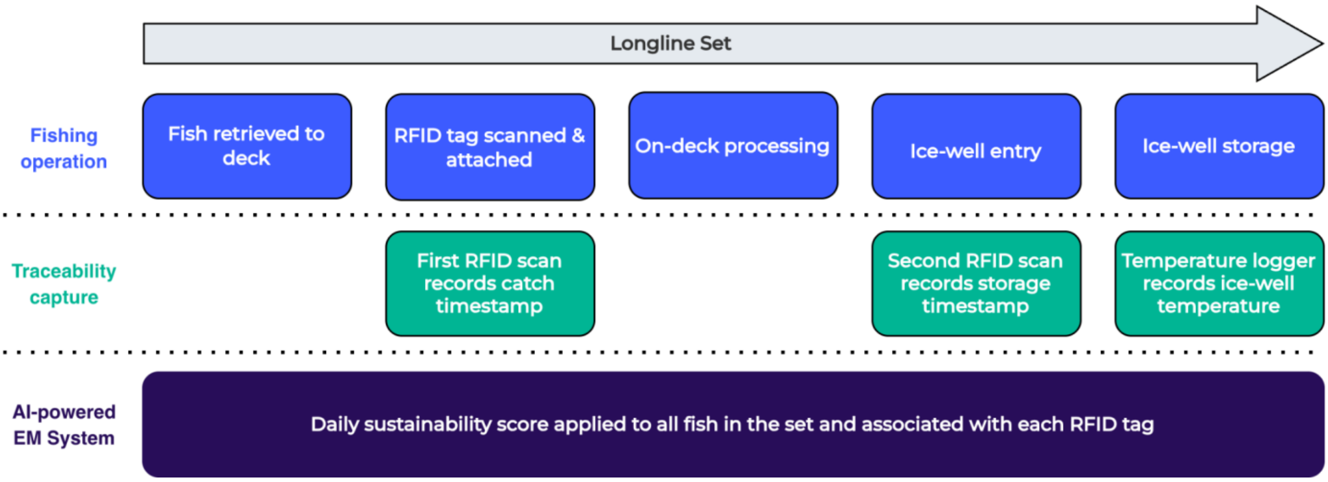

Figure 3: Overview of the traceability workflow. The diagram shows operational fishing steps (top), corresponding traceability data capture events (middle), and AI-generated sustainability scores applied at the set level (bottom).

During a typical longline set, fish are retrieved one by one as the gear is hauled back to the vessel. As each market catch reaches the deck, the crew attaches an RFID tag to the tail (caudal fin). Before attaching it, the tag is scanned by the first RFID antenna, which activates the tag and captures the exact timestamp the fish comes onboard.

After on-deck processing of the catch, the fish is transferred into the ice well for chilling. As it passes through the ice-well hatch, the same tag is scanned again by a second RFID antenna installed at the hatch entry. This scan records the moment the fish enters cold storage. The elapsed time between the first and second scans is automatically calculated as the time-to-ice metric, a key proxy for product quality that reflects how quickly the fish was processed and chilled after capture.

While fish are stored in the ice well, a waterproof temperature logger continuously records the ice-well temperature. For each fish, temperature readings corresponding to its storage interval are summarized as the median well temperature, providing an additional indicator of catch quality.

In parallel, the AI-powered EM system analyzes electronic monitoring footage throughout the fishing day and generates a daily sustainability score. Each fish is later associated with the sustainability score corresponding to the set in which it was caught.

3.2 Hardware setup and installation

Hardware components were selected for their durability in marine environments, compatibility with the existing EM system, and minimal impact on crew workload. The specific devices used in this pilot, along with their configuration and setup instructions, are provided in Appendix B.

3.2.1 CORE COMPONENTS AND PLACEMENT

All hardware components are positioned along the standard processing path of fish during onboard fishing operations (Figure 4).

RFID tags
Each fish is tagged with a unique RFID identifier. Unused tags are stored in a metal box mounted beside the fishing area to prevent unintended reads and to keep the tagging process efficient.

RFID reader and antennas
A single RFID reader supports two traceability scan points. The first antenna is mounted overhead in the fishing area to detect and activate each tag as it is presented during longline retrieval. The second antenna is installed at the ice-well hatch entry to record the moment each fish enters cold storage.

Temperature sensor and gateway
Inside the ice well, a waterproof temperature logger is mounted on the back wall at a height that prevents direct ice contact. A companion gateway located inside the protected equipment rack retrieves temperature logs via Bluetooth and uploads them automatically over the vessel's internet connection, ensuring continuous, hands-off monitoring of well temperature.

Figure 4: Physical installation locations of traceability equipment aboard the partner vessel.

3.2.2 INTEGRATION WITH EXISTING EM HARDWARE

The partner vessel is equipped with a full EM system, including cameras, GPS sensors, onboard servers, satellite internet connectivity, and an NVIDIA Jetson edge device. The traceability hardware was designed to integrate directly into this existing architecture, leveraging shared power, networking, and data resources.

All networked traceability components—including the RFID reader and the temperature gateway—connect via Ethernet to the vessel's protected hardware rack, which houses the Wi-Fi router, network switches, onboard servers, and the edge device. This configuration ensures stable power and data transfer while maintaining a secure local network.

The second RFID antenna installed at the ice well connects to the reader via a coaxial cable, completing the physical link between the fishing area and cold storage. Figure 5 illustrates the physical layout and wiring of the traceability components in relation to the vessel's EM system.

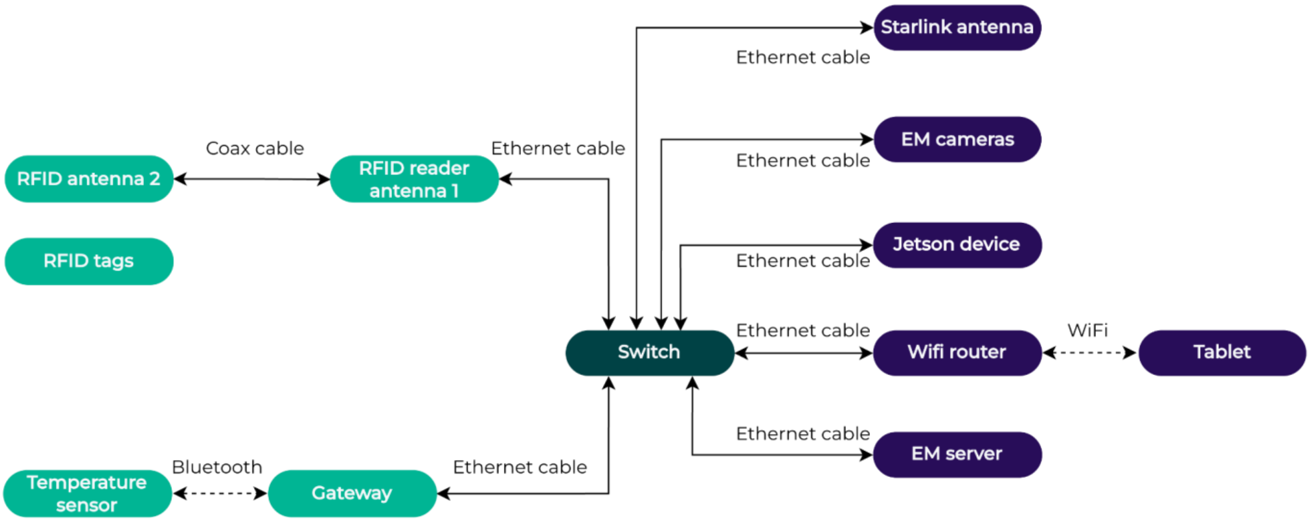

Figure 5: On-vessel hardware wiring diagram, showing physical connections between all installed components.

4. Traceability data

4.1 Full system architecture

Figure 6 illustrates how the traceability system integrates with onboard and cloud-based components. The edge device serves as the central coordinator, bringing together data from the EM system and the traceability hardware described in Section 3.2.

Cameras and GPS sensors from the EM system continuously record fishing activity and vessel position throughout the trip. These data are combined with electronic logbooks submitted by the captain and analyzed by the AI-powered EM system to generate *Daily Reports*, including a sustainability score for each fishing day. The reports are then transmitted to secure cloud storage via the vessel's Starlink satellite connection.

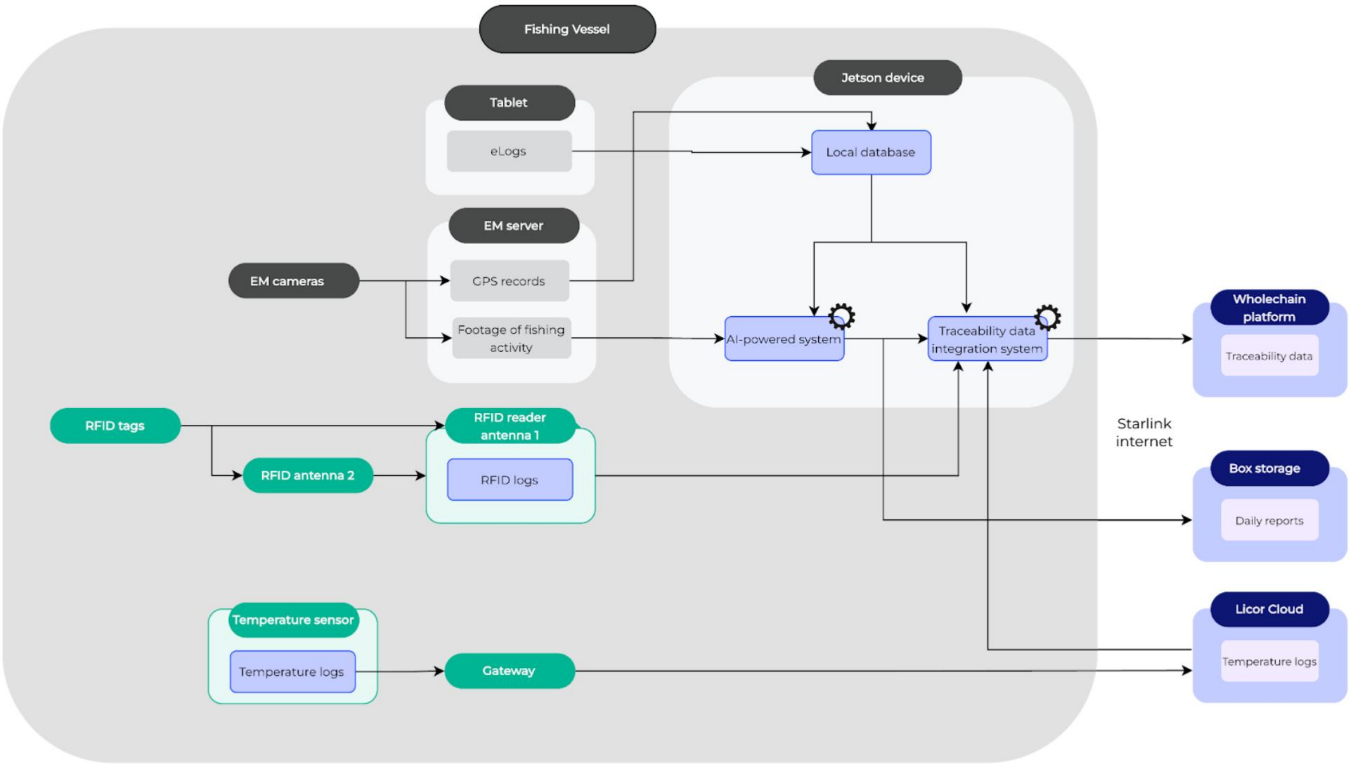

Figure 6: full system architecture diagram showing the interactions between onboard and cloud-based components. The grey area highlights components installed on the vessel, green boxes show hardware and software introduced for first-mile traceability, and blue boxes show cloud services used for data storage, processing, and sharing.

In parallel, fish-level traceability data—such as RFID scan events and ice-well temperature measurements—are captured onboard and processed locally on the edge device. The traceability system integrates these handling and storage data with the AI-generated sustainability scores, creating a complete record for each tagged fish that links catch location, onboard handling, storage conditions, and sustainability conditions (see Section 4.4 for details).

The following sections describe in more detail how these data are captured, stored, and integrated to produce complete first-mile traceability records.

4.2 Data sources

4.2.1 RFID TAG LOGS

RFID tags provide the backbone of the traceability system by giving each fish a unique identifier. Whenever a tag is scanned, the RFID reader generates a log containing the tag ID, the antenna that detected it, and the exact timestamp of the event. These logs are sent directly to the edge device through HTTP requests.

Example RFID log:

```
{
 "epc": "A00000000000000000000127",
 "antenna": 2,
 "timestamp": "2025-12-02T08:32:12Z"
}
```

4.2.2 TEMPERATURE DATA

The temperature sensor records the ice well’s temperature every 30 minutes and sends the data to the cloud through the gateway. During data integration, the edge device retrieves these logs via API and uses them to determine whether each tagged fish was kept within the required temperature range while onboard.

Example temperature record:

```
{
 "logger_sn": "22209760",
 "sensor_sn": "22209760-1",
 "timestamp": "2025-12-02 08:30:00Z",
 "value": 4.03,
 "unit": "°C"
}
```

4.2.3 GPS POSITIONS

The existing EM system continuously records vessel GPS position, which is stored on the internal server and synchronized to a local database in the edge device every 15 minutes. The traceability integration system later cross-references RFID timestamps with these GPS logs to determine the precise location where each fish was caught.

Example GPS log:

```
gps_datetime                  |  lat        |  lon
2025-12-02 09:09:56Z    |7.78188        |-83.6467996
2025-12-02 09:09:50Z    |7.7819235      |-83.6467474
```

4.2.4 SUSTAINABILITY SCORES

At the end of each fishing day, the AI-powered EM system generates a Daily Report that includes a sustainability score. This score summarizes three components: GPS risk, eLog risk, and illegal and/or sensitive species risk.

The edge device stores the Daily Reports locally, and the traceability data integration system retrieves the sustainability score from the corresponding folder when processing the trip's data and associates this score with the relevant tagged fish.

Example risk-score record:

```
{
 "date": "2025-12-02",
 "sustainability_score": 2,
 "components": {
  "elog": 1,
  "illegal_species": 2,
  "gps": 2
 }
}
```

4.2.5 FINALIZE TRIP UI

The edge device hosts a simple user interface (UI) that the captain can access from the tablet to end a fishing trip. This action signals that fishing activity has concluded and that the trip's data can be closed out as a single, complete unit.

When the captain selects "Finalize trip", the traceability system records that timestamp as the official trip close and triggers the integration pipeline, which consolidates all onboard records into a complete traceability dataset, as described in Section 4.4.

```
{
 "event": "finalize_trip",
 "timestamp": "2025-12-02T16:45:10-06:00"
}
```

4.3 Data storage

All traceability data generated onboard flow into a structured onboard relational database, organized across three core tables (Figure 7):

- fish: one row per tagged fish
- trips: one row per fishing trip
- sets: one row per fishing set (derived after the trip)

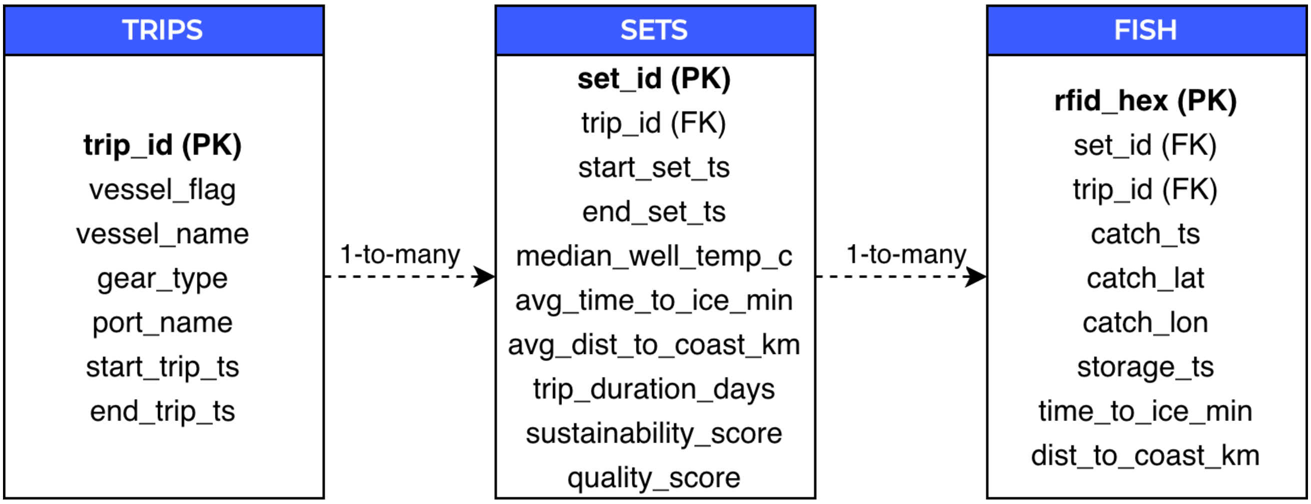

Figure 7: Conceptual traceability data model.

Note: In the implemented database, sets are identified using a composite key (trip_id, set_number), and fish are associated with sets using those fields rather than a standalone set_id.

4.3.1 FISH TABLE

A new fish record is created the moment an RFID tag is scanned for the first time. As the trip continues, that same record is enriched with data from other sensors and processes. This table stores**:**

- The fish's unique RFID identifier
- Catch timestamp from the moment an RFID tag is first scanned (antenna 1)
- Catch location, expressed as latitude and longitude (decimal degrees), derived from the nearest GPS record in time
- Storage timestamp, recorded when the fish enters the ice well (antenna 2)
- Time-to-ice (minutes), calculated as the difference between catch and storage timestamps
- Distance to coast (kilometers), calculated from the catch location
- Set number, linking the fish to its corresponding set
- Trip ID, linking the fish to its corresponding fishing trip

4.3.2 TRIPS TABLE

A trip record is created when the captain selects "Finalize trip" on the tablet as the vessel approaches port. This table stores:

- Vessel and gear information
- The start and end timestamps of the trip
- A unique trip_id used to group all fish caught between those timestamps

Traceability data are shared with Wholechain on a per-trip basis. Each export contains one trip record along with all associated sets and fish.

4.3.3 SETS TABLE

Unlike trips and fish, longline sets are not recorded explicitly. They are inferred after the trip ends by analyzing the sequence of catch timestamps. When the time gap between two consecutive fish exceeds a defined threshold, the traceability system treats this as the boundary between sets and increments the set counter (Figure 8). Once these boundaries are established, each set is summarized using the fish assigned to it, along with temperature data and sustainability scores. This table contains:

- Start and end timestamps of the set
- Median ice-well temperature (in Celsius) during the set's storage window
- Average time-to-ice (in minutes) across all tagged fish in the set
- Average distance to coast (in kilometers) across all tagged fish in the set
- Trip duration
- A traffic-light sustainability score (aggregated from daily AI risk scores)
- A traffic-light quality score (derived from temperature + time-to-ice + trip duration)

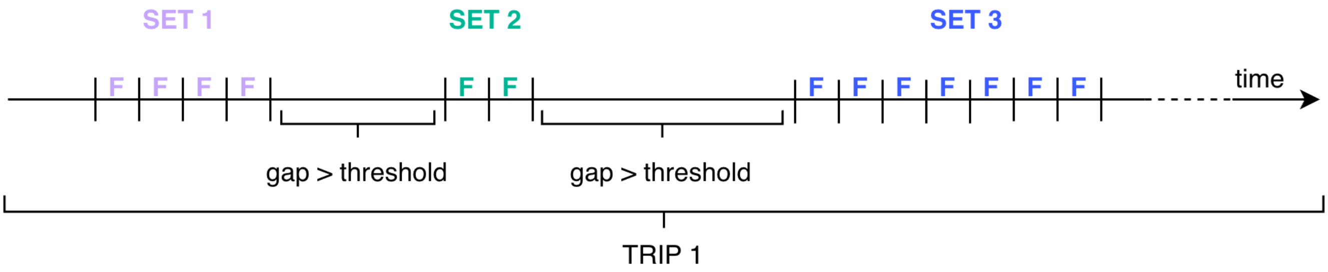

Figure 8: How longline sets are inferred from catch timestamps. Each "F" represents a tagged fish along the timeline of a single trip. Long gaps between catches are used to identify set boundaries.

4.4 Data integration system

All traceability data collected onboard are consolidated by a dedicated software module running on the vessel's edge device. This traceability integration system acts as the coordinating layer that connects the data sources described in Section 4.2 with the storage model outlined in Section 4.3, transforming raw sensor inputs and AI outputs into complete, structured traceability records.

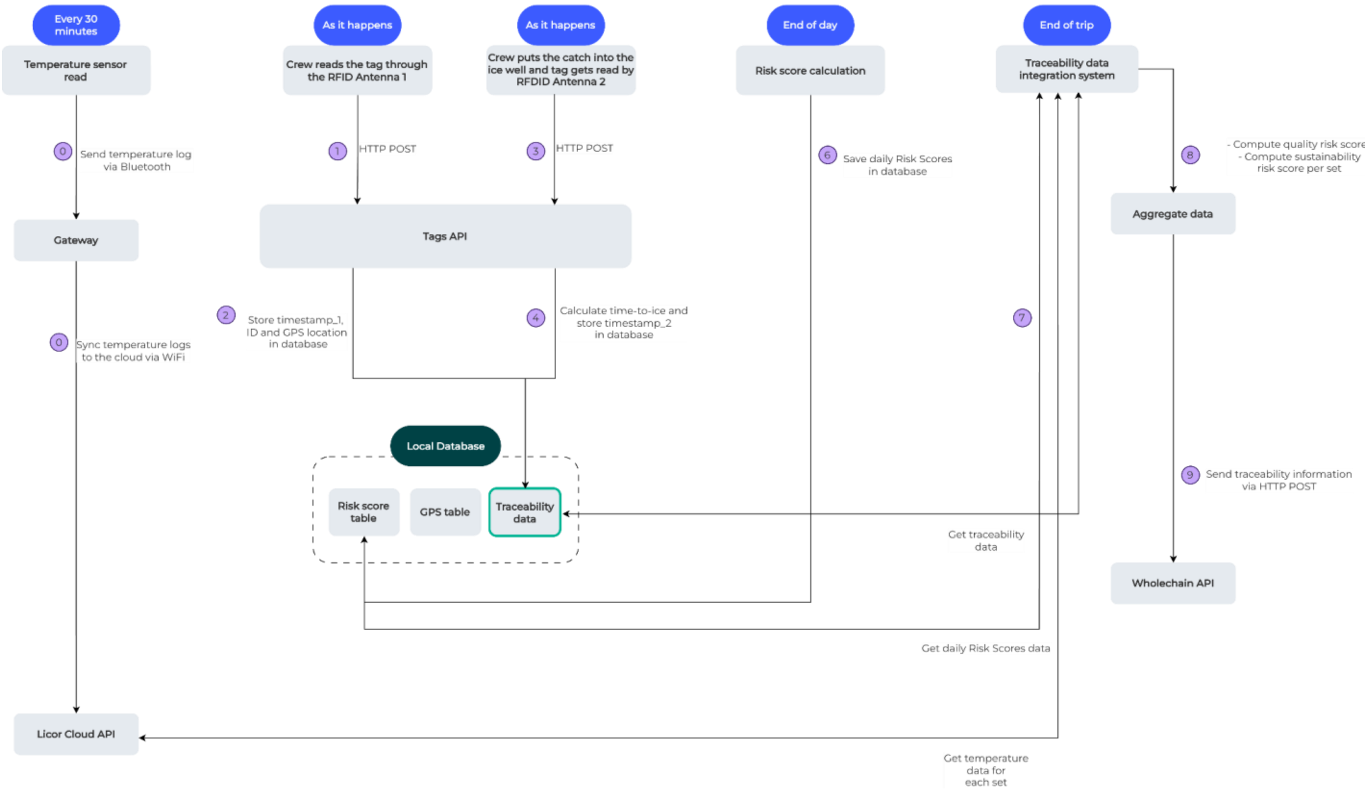

Figure 9: diagram showing how data flows between the different software components inside the vessel and on the cloud. The numbers describe the sequence of events, and the blue markers highlight the exact times each event is triggered.

As illustrated in Figure 9, the integration system operates in two phases. During the trip, it incrementally enriches fish records as data become available. RFID scan events create and update fish entries, which are spatially anchored using the nearest GPS position in time. Subsequent scans at the ice-well hatch complete each fish's handling timeline by recording storage entry and enabling calculation of time-to-ice. In parallel, temperature measurements from the ice well and daily sustainability scores generated by the AI-powered EM system are stored locally and associated with the relevant fishing periods.

The full interpretation of these data streams occurs once the captain finalizes the trip. At this point, the integration system closes the trip record, assigns a single trip ID to all fish caught within the trip window, and infers longline sets by analyzing gaps between successive catch timestamps. This step converts the chronological flow of fish into operational units that reflect how the crew actually fished.

Once sets are identified, the traceability system summarizes both quality and sustainability information at the longline set level. For quality, the traceability system uses three proxies that capture storage conditions and handling practices: the median well temperature during the storage period of the set, the average time-to-ice across all fish in the set, and the overall trip duration. The current implementation applies the following initial thresholds:

Table 2. Quality proxies and associated numerical thresholds used to calculate the quality score.

Quality Proxy	Description	Green	Yellow	Red
Well temp	How consistently the cold chain was maintained	≤ 4.4 °C	≤ 10 °C	> 10 °C
Time-to-ice	How quickly the fish was sacrificed, de-bled, and chilled	≤ 90 min	≤ 180 min	> 180 min
Distance to coast	Distance from shore at the time of capture	> 50 km	≥ 25 km	≤ 25 km
Trip duration	How long fish remained onboard before landing	≤ 7 days	≤ 21 days	> 21 days

These thresholds are fully tunable and were selected as reasonable starting points for the pilot. They reflect values suggested by practitioners, but they are not intended to represent universal or final standards. As more data are collected, they can be adjusted to better match real quality outcomes. In the color-coded quality scale, green represents ideal or high-quality handling, yellow indicates acceptable conditions with some room for improvement, and red signals practices and conditions likely to reduce seafood quality.

Sustainability scores are summarized similarly at the set level. For each set, the traceability system aggregates the AI-generated daily sustainability scores associated with the fishing day(s) spanned by that set, producing a single set-level sustainability score. This score is likewise mapped to a traffic-light scale, where green indicates low risk, yellow reflects moderate concerns, and red signals higher-risk IUU-related activity such as fishing near protected areas, potential underreporting of catch, or AI detections of illegal or sensitive species.

With these elements in place, the traceability system compiles every fish, set, and trip attribute into a unified and relational dataset. The structure and transmission of this dataset to downstream systems are described in Section 4.5.

4.5 Data output sent to Wholechain

To ensure that traceability data collected on the vessel can be reliably shared and interpreted downstream, the traceability system exports information in a format aligned with the Global Dialogue on Seafood Traceability (GDST). GDST defines a common data language for seafood traceability, allowing information to move consistently from harvest through processing, distribution, and ultimately to buyers and regulators. Using a GDST-aligned structure ensures that the data generated onboard is compatible with existing buyer systems and regulatory requirements and avoids the need for custom, one-off integrations.

While this pilot partnered with Wholechain as the buyer-facing platform, the traceability system itself is modular. The same finalized dataset can be mapped to any platform or service that supports GDST-compliant traceability events, allowing flexibility as buyer needs and technology partners evolve.

Once a trip is finalized, the traceability integration system assembles all fish-, set-, and trip-level information into a single structured payload.

At a high level, the payload contains:

- One fishing (commission) event per fish, linked to a unique RFID identifier. Each event includes the time and location of catch, along with key traceability attributes such as set number, time-to-ice, distance to coast, quality score, and sustainability score.
- A single shipping event for the trip, representing the transfer of all fish from the vessel to the next supply-chain location (e.g., port or buyer). This event groups all individual fish into a coherent shipment and marks the transition from onboard handling to downstream logistics.

A simplified representation is shown below to illustrate the structure and content (field names and values are indicative):

```
{
 "events": [
  {
   "type": "commission",
   "event_time": "2025-11-14T05:43:49Z",
   "location": {
    "vessel_name": "Anonymized Vessel",
    "vessel_flag": "CR"
   },
   "product_instance": {
    "rfid": "A00000000000000000000127",
    "product_id": "TUNA"
   },
   "catch_location": {
    "latitude": 8.279186,
    "longitude": -83.6145205
   },
   "custom_properties": {
    "avg_time_to_ice_min": 159,
          "median_well_temp_c": 7.2,
    "avg_dist_to_coast_km": 18.3,
    "trip_duration_days": 6.9 ,
    "set_number": 1,
    "trip_id": 1,
    "sustainability_score": "Yellow",
    "quality_score": "Yellow"
   }
  },
  {
   "type": "ship",
   "event_time": "2025-11-21T02:36:00Z",
   "ship_from": "Fishing Vessel",
   "ship_to": "Buyer Location",
   "product_instances": [
    { "rfid": "A00000000000000000000127" },
    { "rfid": "A00000000000000000000147" }
   ]
  }
 ]
}
```

Together, these events form a complete GDST-aligned record that preserves fish-level granularity while remaining compatible with shipment-based supply-chain systems.

After the payload is constructed, it is transmitted automatically to Wholechain via a secure API connection. Built-in validation and retry mechanisms ensure that data are delivered reliably, even under intermittent connectivity conditions. Once received, the data become immediately available within the Wholechain platform for buyer access and verification.

From a buyer's perspective, Wholechain serves as a window into these traceability data as products move through the supply chain. Buyers can view shipments received from vessels, see which individual fish are included, and explore fish-level details linked back to the original fishing event. For each fish, buyers can access catch timing and location, trip and set information, and the associated quality and sustainability scores.

5. Software development

This section outlines the key practices, tools, and workflows used to ensure robust, maintainable, and high-quality software throughout the project lifecycle.

5.1 Codebase

The source code is publicly available under TNC's GitHub repository. The codebase uses a structured Git workflow (Figure 10) to ensure a clear separation between active development and stable releases. The *main* branch serves as the production-ready codebase deployed to the vessel's edge device. New features were developed in isolated branches created from *develop*, which acted as the integration branch. Once reviewed and validated, *feature* branches were merged into *develop*, and when a release was ready, *develop* was merged into *main*.

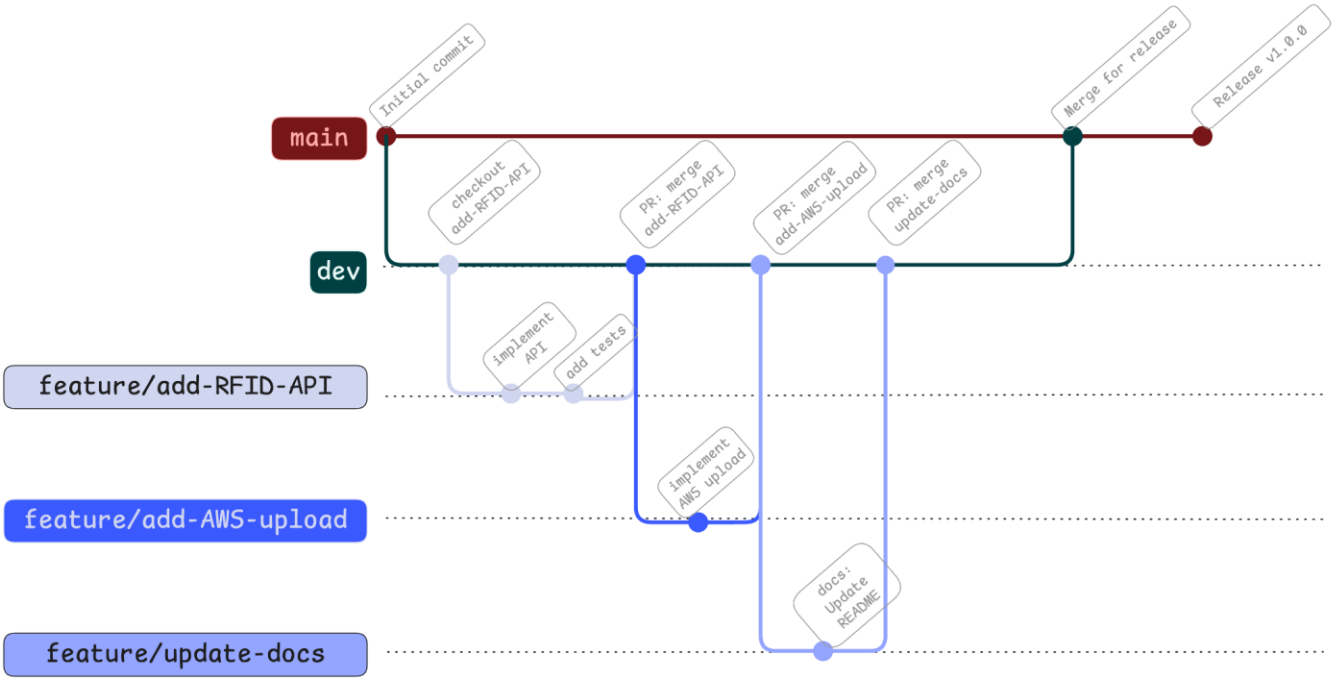

Figure 10. The code lifecycle from development to a stable release.

The package manager *uv* was used to ensure efficient dependency management and reproducible environments, while *Ruff* was used to enforce code style consistency and perform fast, reliable static code analysis across the codebase.

5.2 Testing

Code quality was safeguarded through a comprehensive automated testing framework, triggered via pre-commit hooks and GitHub Actions as part of Continuous Integration (CI). This process automatically ran checks on every proposed change before merging, including:

- Static analysis with Ruff to enforce consistent style and detect early issues
- Secret scanning to prevent accidental exposure of sensitive credentials
- Execution of a unit test suite covering core components such as metric calculations, data aggregation, and tracking utilities

In addition to automated CI checks, candidate releases were tested in a dedicated staging environment that mirrored the vessel hardware. An edge device on Tryolabs' premises was configured identically to the deployed AI-powered system. All code was run end-to-end under real-world constraints to ensure performance and stability before deployment.

5.3 Reviewing

In addition to automated checks, all new code contributions required a mandatory peer code review. At least two other developers reviewed each submission to ensure logical correctness, adherence to architectural patterns, and overall code quality before merging it into the development branch.

Leveraging this template-driven approach streamlined the development process, reduced setup time, and reinforced best practices throughout the project lifecycle.

6. Pilot deployment and results

The pilot deployment comprised two commercial longline fishing trips conducted on the same vessel operating in the Eastern Tropical Pacific. Across both trips, the crew tagged a total of 39 tunas, enabling the traceability system to generate complete fish-level records and GDST-aligned trip payloads. These datasets were successfully transmitted to Wholechain, where downstream supply chain actors can now explore the sustainability and quality scores of each fishing set and tagged fish, respectively.

From an operational standpoint, the on-vessel workflow integrated smoothly into routine fishing activities. The captain reported that the tagging process was easy for the crew to adopt and did not interfere with standard handling practices, demonstrating that high-resolution traceability can be implemented with minimal operational burden.

6.1 Traceability outputs generated during the pilot

Across the two trips, tuna were tagged during six longline sets. For each set, the traceability system aggregated fish-level RFID events, sensor readings, and AI-generated sustainability scores into a set-level summary that reflects both onboard handling conditions and sustainability conditions. Tables 3.a. and 3.b. present anonymized summaries of the results from the first and second trips, respectively.

Table 3.a. Set-level summary from pilot's 1st trip (anonymized)

set_number	no_tagged_tuna	median_well_temp_c	avg_time_to_ice_min	avg_dist_to_coast_km	trip_duration_days	sustainability_score	quality_score
1	8	7.5	159	18	6.9	Yellow	Yellow
2	13	7.4	109	38	6.9	Green	Yellow
3	5	6.9	137	38	6.9	NA	Yellow
4	7	6.7	156	58	6.9	NA	Yellow

Table 3.b. Set-level summary from pilot's 2nd trip (anonymized)

set_number	no_tagged_tuna	median_well_temp_c	avg_time_to_ice_min	avg_dist_to_coast_km	trip_duration_days	sustainability_score	quality_score
1	1	5.9	40	31	13.3	Yellow	Yellow
2	5	6.5	121	53	13.3	Yellow	Yellow

6.1.2 QUALITY SCORE INTERPRETATION

Using the thresholds defined in Table 2, all sets in both trips received a yellow quality score, but for different underlying reasons in each trip.

During the first trip (Table 3.a), handling conditions were relatively stable across all four sets. Median well temperatures consistently fell between 6.7 and 7.5 °C, and average time-to-ice ranged from 109 to 159 minutes—both within the yellow threshold. Trip duration remained short (6.9 days), placing it in the green range. When these indicators were combined—two yellow proxies (temperature and time-to-ice) and one green proxy (trip duration)—the resulting quality score for every set was yellow.

In the second trip (Table 3.b), handling conditions differed in magnitude but led to the same overall classification. Median well temperatures remained within the yellow range, while time-to-ice varied more noticeably between sets—from 40 minutes (green) in Set 1 to 121 minutes (yellow) in Set 2. However, the longer trip duration (13.3 days) shifted from green to yellow, offsetting improvements in time-to-ice for some fish. As a result, the combined indicators again produced a yellow quality score for both sets.

6.1.3 SUSTAINABILITY SCORE INTERPRETATION

Sustainability scores were generated for all sets where complete AI monitoring data were available.

During the first trip (Table 3.a), Set 1 received a yellow sustainability score driven by AI detections of retained species subject to regulatory protections under certain size thresholds in the region where the vessel operates. Set 2 received a green sustainability score, indicating no elevated sustainability concerns during that fishing activity. Sets 3 and 4 are marked as “NA” because temporary hardware issues limited complete data capture for those periods.

In the second trip (Table 3.b), both Sets 1 and 2 received yellow ratings. As in the first trip, these yellow scores were driven by detections of retained regulated species. This result illustrates how the traceability system can surface concrete, actionable signals—such as interactions with regulated species—that may warrant follow-up, crew training, or adjustments to handling practices.

6.2 Key findings from the pilot

Overall, the pilot demonstrated several important outcomes:

- **End-to-end first-mile traceability is feasible in commercial settings**, producing fish-level records that link catch location, handling practices, and sustainability scores.
- **AI-powered sustainability scores can be operationalized for markets**, moving beyond compliance review into buyer-facing signals.
- **Product quality proxies can be derived from onboard sensor data**, enabling characterization of first-mile handling and storage conditions.
- **Complex traceability data can be translated into simple, actionable signals**, allowing quick assessment while preserving access to underlying data.
- **The system integrates with existing traceability platfor**ms, with onboard data successfully transmitted to Wholechain.
- **Crew adoption was strong**, showing that high-resolution data collection does not require disruptive changes to routine operations.

6.3 Technical learnings and areas for refinement

The pilot revealed several areas for future innovation and refinement. The RFID antenna installed at the ice-well hatch performed reliably under dry conditions but experienced reduced reading accuracy when exposed to splashes. A thin layer of water on the antenna surface was enough to interfere with tag detection, and although wiping the surface resolved the issue, this is impractical during active fishing. Future iterations may incorporate a splash shield, an alternative mounting angle, or hardware specifically designed for wet, high-impact environments.

The pilot highlighted the importance of minimizing the physical footprint of onboard hardware. Feedback from the vessel indicated that equipment size and placement can affect crew operations, underscoring the need for leaner designs. Future deployments should prioritize compact, IP-rated components that can be distributed across available spaces on the vessel, rather than relying on a single centralized hardware rack.

The pilot also showed that the initial quality metrics—selected as the most readily measurable indicators in this pilot—capture only part of the story. Time-to-ice, for instance, is not a definitive indicator of quality: fish chilled too quickly while still warm can burn, while fish kept shaded and periodically wetted may maintain good quality before icing. Similarly, the well-temperature sensor measures air temperature, not the temperature of the fish in the ice. In tropical fisheries, opening the well temporarily increases air temperature even though the fish remain cold. These limitations were expected at this early stage, and the pilot has already prompted discussions with industry experts about additional sensors and more accurate proxies. A key next step is to compare onboard indicators with the grades assigned at landing to test how well the proxies reflect actual quality.

Despite these early-stage limitations, the pilot generated strong interest from processors and buyers who reviewed the outputs. Stakeholders emphasized that having objective, fish-level and set-level information—especially around quality and sustainability—is valuable for purchasing decisions.

7. Conclusions and next steps

The traceability system generates a uniquely granular dataset that integrates sustainability and quality information, marking an important step toward making first-mile traceability practical and scalable in commercial longline fisheries. By showing that high-quality traceability data can be generated onboard and integrated seamlessly into existing supply-chain systems, the pilot demonstrates a clear pathway for improving first-mile transparency. Beyond its technical achievements, the project underscores the role of first-mile traceability in aligning incentives across the supply chain. Making quality and sustainability performance visible at the set level strengthens the connection between onboard practices and downstream outcomes such as product value, buyer preferences, and market access. This transparency enables crews and vessel managers to better understand how fishing and handling decisions influence results, supporting targeted feedback and ultimately on-the-water improvements. At the same time, buyers gain access to objective, shared information, enabling more constructive engagement with suppliers and strengthening trust through consistent, verifiable sourcing.

The findings also clearly define priorities for the next phase of development. Expanding deployment to additional vessels will be essential to test system robustness across a wider range of operating conditions and crew practices. Technical refinements—including more compact and robust hardware, improved sensor placement, and expanded quality proxies—will further enhance system reliability. Equally important, validating onboard quality indicators against quality assessments at port will help ensure that metrics generated on the vessel accurately reflect real product outcomes.

Together, these results illustrate how first-mile traceability can transition from pilot to standard practice. With continued refinement and validation, the technical blueprint presented here provides a scalable foundation for more transparent supply chains, better-aligned incentives, and sustained improvements in both product quality and responsible fishing practices.

8. References

Food and Agriculture Organization (FAO) of the United Nations. (2020). *The state of world fisheries and aquaculture 2020: Sustainability in action.* https://doi.org/10.4060/ca9229en

Tryolabs. (2026). *Mobilizing AI and edge technologies to advance near-real-time electronic monitoring footage review* (Report prepared for The Nature Conservancy) https://www.nature.org/en-us/what-we-do/our-insights/ai-electronic-monitoring-fisheries-report/

Appendix A: training guide for the crew

This appendix provides step-by-step instructions for vessel crew on how to use the onboard traceability equipment during routine fishing operations.

A.1 RFID data collection points

RFID tags are scanned at two moments during routine fishing operations: at catch retrieval and when the fish enters the ice well. No additional steps are required beyond those described below.

CATCH RETRIEVAL

1. Remove the fish from the longline as usual.
2. Take one RFID tag from the metal storage box.
3. Hold the tag under the overhead RFID reader until the orange LED confirms the scan.
4. Attach the tag just below the tail fin using the zip tie. Tighten firmly.

ICE WELL ENTRY

1. After cleaning and handling, pick up the fish to move it into the ice well.
2. As the fish passes through the hatch, the RFID antenna at the well entrance will automatically scan the tag.
3. Check for the orange LED on the reader to confirm a successful scan.

Important note: If either RFID antenna becomes very wet, reading performance may decrease. If this happens, wipe the antenna dry with a cloth before continuing.

Appendix B: hardware selection and configuration

This appendix documents the hardware components used in the pilot traceability system and summarizes their configuration and setup. It is intended as a technical reference for the deployment, installation, and verification of the onboard traceability equipment.

B.1 Hardware selection

Component	Description
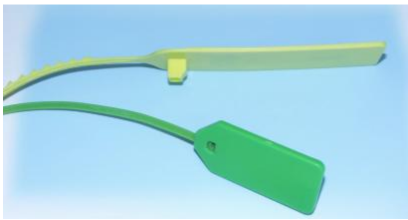	RFID Tags Model: GAO-116092 (Zip Lock Durable Waterproof RFID Tag with QR code)
	RFID Reader (Integrated Antenna) Model: GAO-216031 (UHF RFID IP67 Integrated Reader/Writer with built-in antenna) [Provider's website]
	External RFID Antenna Model: GAO-326005 (UHF 900 MHz 7.5 dBi circular antenna with 10 ft RF cable and mounting bracket) [Provider's website]
	Temperature Sensor Model: HOBO TidbiT MX2203 (400-ft waterproof Bluetooth data logger) [Provider's website]
	Gateway Model: MXGTW1 (MX Gateway) [Provider's website]

B.2 Hardware configuration

This section describes the configuration and verification steps for the onboard hardware used in the pilot. The procedures below were followed during installation and are included to support replication or future deployments.

B.2.1 GATEWAY AND TEMPERATURE SENSOR

Device registration and account setup

- Register both the gateway and the temperature logger in LicorCloud.

1. Connect the gateway to power using the AC adapter.

- Download the HOBOconnect mobile application (iOS or Android).
- Open the app and link it to the LicorCloud account *(Settings → Connect Account).*

Gateway configuration

1. In Devices, wait for the gateway to appear.
2. Select the gateway and choose Configure & Start.
3. Assign a recognizable name (e.g., Vessel Gateway).
4. Configure network connectivity:
    - Ethernet (DHCP) for wired internet, or Wi-Fi, entering the network name and password.
5. Tap Start to save settings.

Verify

- Gateway LED shows a solid green light.
- In HOBOconnect, the gateway status reads “Gateway running”.

Logger configuration

1. In Devices, select the temperature logger and choose Configure & Start.
2. Assign a name (*e.g., Ice Well Logger*).
3. Set Upload Data Via to Gateway.
4. Configure:
    - Logging interval: 30 minutes
    - Start: Immediate
5. Tap Start to save.

Verify

- In HOBOconnect, the logger status reads “Logging”.
- Both devices appear in LicorCloud → Devices → MX Devices.

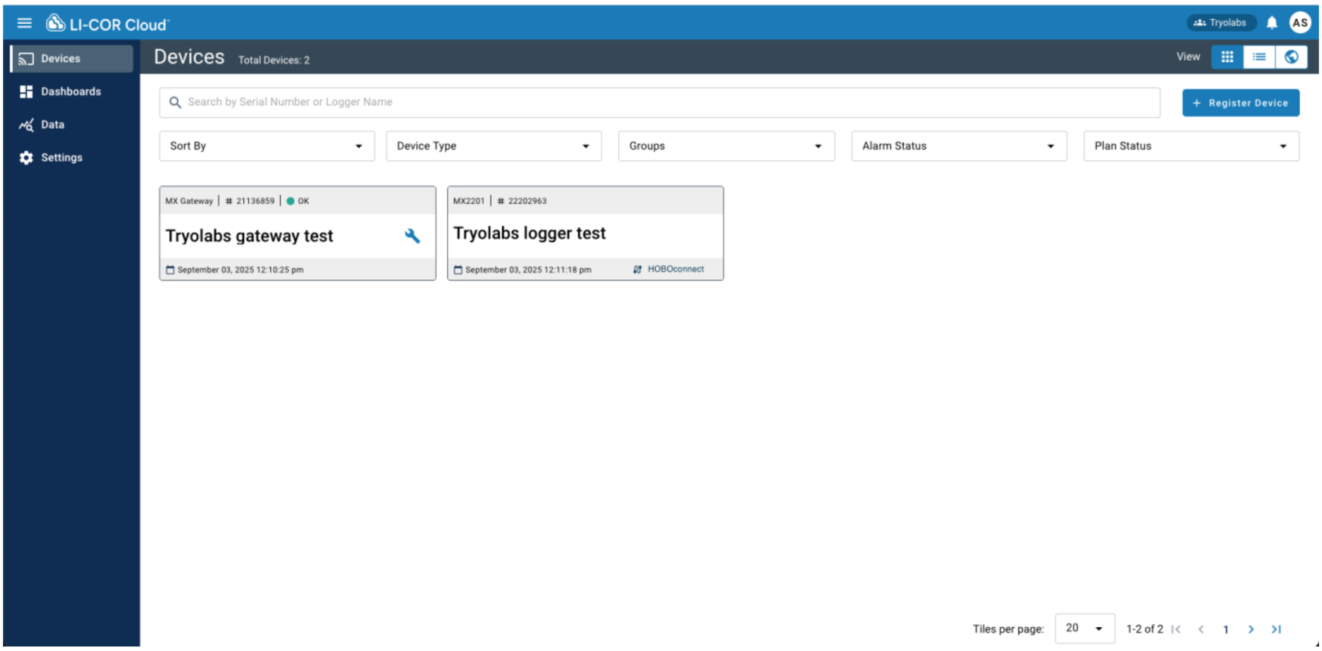

Figure B.1. LicorCloud device list showing the registered gateway and temperature logger under Devices.

Data validation

1. **Allow 15–20 minutes for the first readings.**

- In LicorCloud → Data, select the logger and confirm temperature values are visible and time-aligned.

Bluetooth range testing

1. **Fix the gateway in its intended onboard location.**
2. **Move the logger away in stages:**
    - **Start at 1 m, then increase in 5 m increments (up to ~30 m).**

- At each distance, verify data reception in LicorCloud ⟶ Devices ⟶ Last data.

3. **Repeat with walls or doors between devices.**

Verify

- Line-of-sight range ≈ 30 m
- **Reduced range with obstructions (typically 10–15 m)**
- **Identify a safe operating distance for the vessel layout.**

Power recovery test

1. **Disconnect gateway power for at least 10 minutes.**

- **Reconnect power and observe LED behavior:**
    - Solid yellow → flashing yellow (4–5 minutes)
- **Confirm gateway reappears in LicorCloud as Running.**

2. **Verify at least one new temperature reading uploads after restart.**

B.2.2 RFID READER AND ANTENNA 2

Hardware connection

1. Connect the reader to power using the supplied adapter.
2. Allow 30–60 seconds for startup.
3. Connect the external antenna:
    - Coaxial cable (TNC end) → reader ANT port
    - Other end → external antenna

Verify LEDs

- PWR: solid (power present)
- ANT1: solid (antenna 1 detected)
- ANT2: solid (antenna 2 detected)
- Other LEDs blink according to network status.

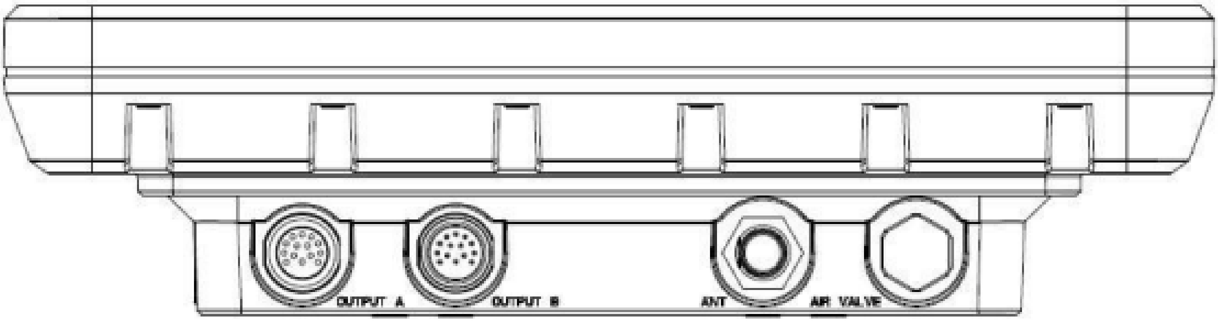

Figure B.2. Physical connections for the RFID reader, including power supply, network connection to the router, and external antenna cabling.

Network connection

1. **Connect the reader to the onboard router using an RJ45 cable**
   (Output A → router LAN port).
2. **Wait ~1 minute for DHCP assignment.**
- Use RFIDReaderTool.exe from the Demo → *Connect → Search* **to identify the reader IP.**

3. **Record the assigned IP address.**

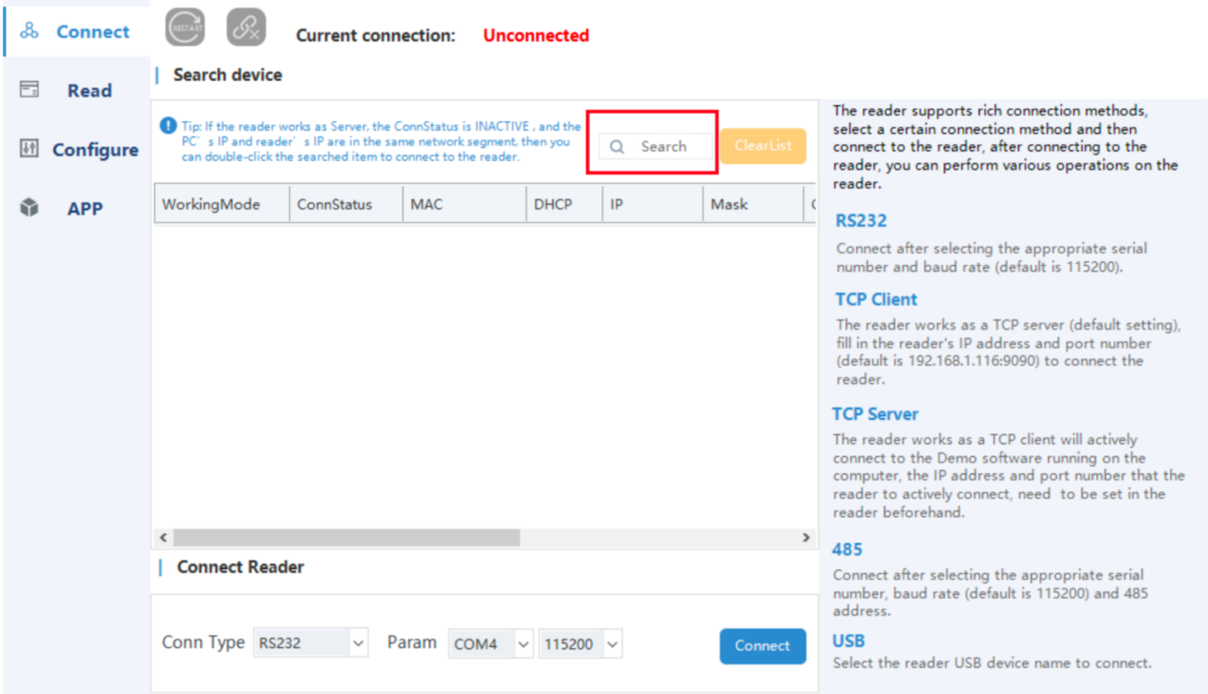

Current connection: Unconnected

Tip: If the reader works as Server, the ConnStatus is INACTIVE , and the PC's IP and reader's IP are in the same network segment, then you can double-click the searched item to connect to the reader.

Conn Type: RS232 Param: COM4 115200

The reader supports rich connection methods, select a certain connection method and then connect to the reader, after connecting to the reader, you can perform various operations on the reader.

RS232

Connect after selecting the appropriate serial number and baud rate (default is 115200).

TCP Client

The reader works as a TCP server (default setting), fill in the reader's IP address and port number (default is 192.168.1.116:9090) to connect the reader.

TCP Server

The reader works as a TCP client will actively connect to the Demo software running on the computer, the IP address and port number that the reader to actively connect, need to be set in the reader beforehand.

485

Connect after selecting the appropriate serial number, baud rate (default is 115200) and 485 address.

USB

Select the reader USB device name to connect.

Figure B.3. RFID ReaderTool demo interface used to discover the reader on the local network (Demo → Connect → Search) and identify the assigned IP address.

Access RFID Manager System

- Open a browser and navigate to: http://<READER_IP>:8080

1. Log in with default credentials: admin / admin.

Verify

- RFID Manager dashboard loads successfully.

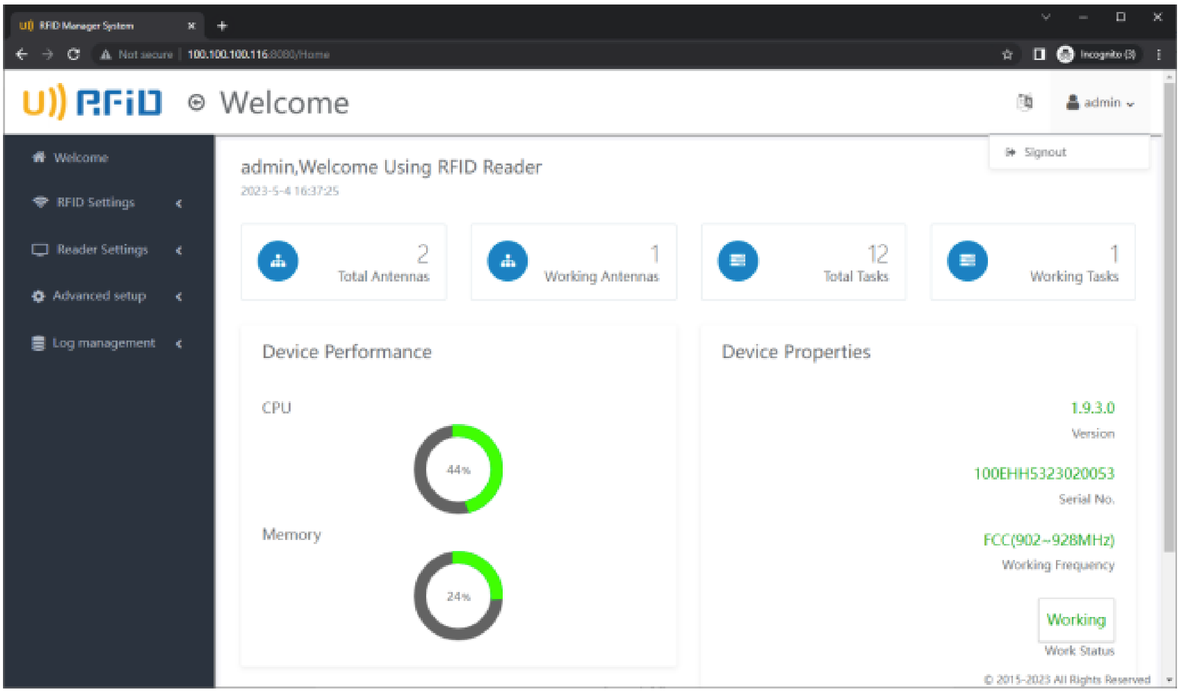

100.100.100.116:8080/Home

admin,Welcome Using RFID Reader

2023-5-4 16:37:25

2 Total Antennas

1 Working Antennas

12 Total Tasks

1 Working Tasks

Version: 1.9.3.0

Serial No.: 100EHH5323020053

Working Frequency: FCC(902~928MHz)

Work Status: Working

Figure B.4. RFID Manager System web dashboard, used to configure antenna settings, reader parameters, and review RFID tag logs (in test mode).

Reader configuration

1. Set Work Status to Idle.
2. Navigate to RFID Settings → Antenna Power:
    - Enable ANT1 and ANT2
    - Set power levels:
        - ANT1: low power (e.g., 2 dBm)
        - ANT2: higher power (e.g., 16 dBm)
3. Configure reader filters:
    - RepeatTime: 5000 ms (prevents duplicate reads)
4. Under Output Settings:
    - Enable Auto Read
    - Enable Local Tag Log (testing) or configure HTTP POST Upload.
5. Set Work Status back to Working.

Functional testing

1. Present a single RFID tag to each antenna.
2. Confirm:
    - Orange LED flashes on the reader
    - Audible beep (if enabled)
    - EPC, timestamp, and antenna ID appear in Log Management → Tag Log
3. Verify timestamps and time zone accuracy.

Load testing (Antenna 2)

1. Present 5–10 tags simultaneously to Antenna 2.
2. Confirm:
    - Multiple EPCs appear in logs
    - No missed reads or system instability

Power recovery test

1. Confirm normal reads with a test tag.
2. Power off the reader for ~30 seconds.
3. Restore power and wait 30–60 seconds.
4. Present a tag again.
5. Confirm:
    - Reader resumes operation without reconfiguration.
    - New reads appear in the Tag Log.

PREPARED BY

tryo·labs

PREPARED FOR

The Nature Conservancy