	Document	Technical White Paper
	Subject	Subsea 3D Imaging for Asset Integrity Teams
	Service	Advanced Unmanned Robotics Services (AURS)

FROM 2D VIDEO TO SURVEY-GRADE 3D

A Practical Guide to Subsea Digital Twins for Asset Integrity Teams

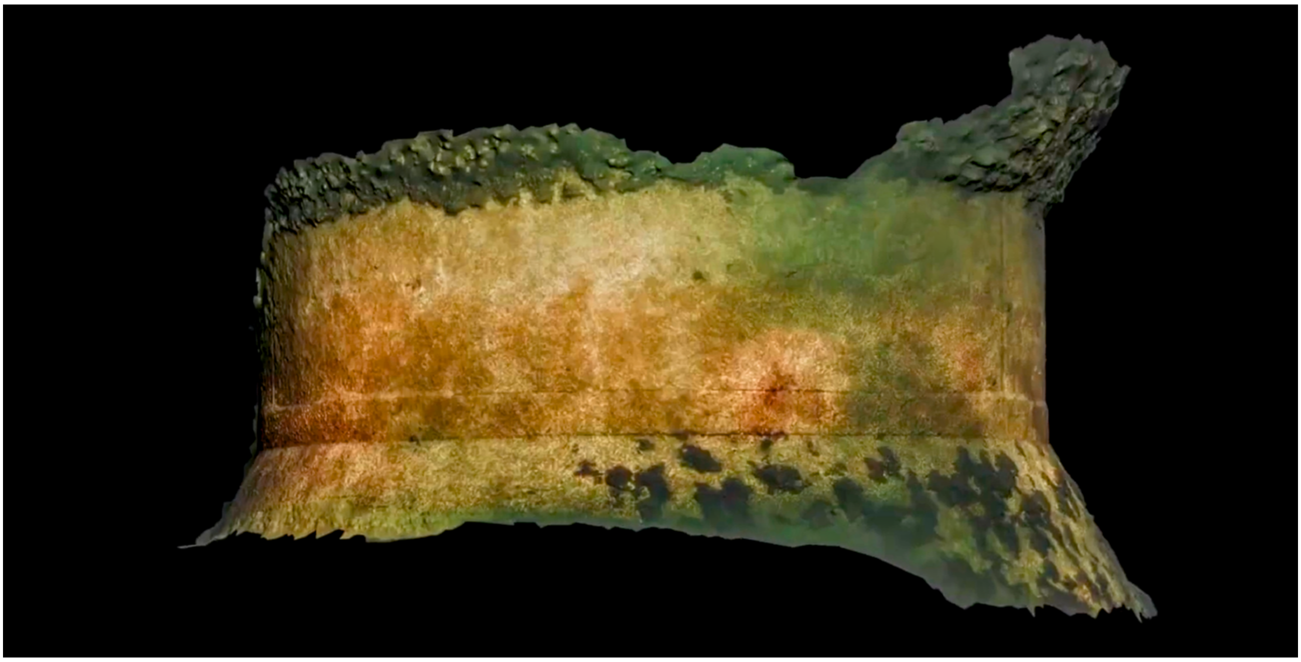

Audience: Asset integrity engineers, inspection planners and ROV/marine coordination teams — Ports, Offshore Wind, Offshore Oil & Gas

Purpose: A practical, vendor-accountable reference for evaluating and adopting subsea 3D digital imaging

Published by: DroneQ Robotics B.V.

Date: August 2026

Classification: Public

DroneQ Robotics B.V.
Slotenmakerstraat 3
6101SK Echt (L), The Netherlands
info@droneq.nl
www.droneq.nl

CONTENTS

1. EXECUTIVE SUMMARY

Most offshore and marine asset integrity programmes still run on two inspection methods that have barely changed in twenty years: a diver or ROV pilot narrating what they see into a 2D video feed, and a written report built afterward from memory, timestamps and still frames. Both methods work, in the sense that they have kept assets safe for decades. Neither method produces a measurement. Neither produces something you can compare, pixel for pixel, to the same structure a year later. And neither produces something a junior engineer can review without watching hours of footage.

Survey-grade 3D photogrammetry — sub-millimetre-accuracy, georeferenced 3D models built from stereo or multi-view camera data — solves the measurement problem. It has existed in usable offshore form for several years now. Adoption, however, has been slower and messier than the technology's capability would suggest, for reasons that have very little to do with the technology itself: unclear accuracy claims, data formats that don't talk to existing systems, no plan for what to do with a model once it's captured, and no one on the client side trained to use it.

This guide is written for asset integrity engineers, ROV/marine coordination leads and inspection planners who are being asked — by insurers, by class societies, by their own management — to move toward 3D-based subsea inspection, and who want a plain, non-promotional account of what it actually involves before they commit a budget line to it. It is deliberately light on marketing language and heavy on the practical questions that determine whether a 3D imaging programme succeeds or quietly gets shelved after the first pilot.

> *This is a practitioner's guide, not a sales brochure. Where a claim needs a caveat, we've included the caveat. Where 3D imaging is the wrong tool for the job, we've said so.*

2. WHY THIS GUIDE EXISTS

If you manage inspection programs for offshore wind foundations, platform jackets, pipelines, quay walls or port infrastructure, you already know the limitations of the status quo, because you work around them every day:

- Subjectivity: a defect's severity depends on which pilot or diver is describing it, and how they were feeling that day. Two inspections of the same weld, six months apart, are not directly comparable unless the same person did both and remembers what they saw the first time.
- No baseline for change detection: without a measured 3D reference, "has this scour pit gotten worse" is answered by comparing two videos side by side and guessing, not by measuring.
- Report review is a bottleneck: a 40-minute dive video becomes a report someone has to watch in close to real time to QA properly. In practice, most footage is watched once, if at all, and archived.
- Data has a shelf life of one report: once the report is issued, the underlying video is rarely revisited. There is no persistent, queryable model of the asset that the next inspection, or the next contractor, can build on.
- Diver exposure and cost: where divers are still used, every hour underwater is a safety exposure and a day-rate cost that 3D-capable ROV platforms can often reduce or remove entirely.

None of this means 2D video and diver reports are worthless — for a quick visual check, a known-good asset, or an emergency response, they remain fast and adequate. The problem is that they are being used, by default, for integrity-critical, repeat-cycle inspection where a measured, comparable, persistent record would materially reduce risk and cost over the asset's life. That is the specific gap this guide addresses.

3. WHAT "SURVEY-GRADE" ACTUALLY MEANS

"3D scanning" is used loosely enough in this industry that it is worth being precise about what survey-grade actually requires, because the difference between a nice-looking 3D model and a usable engineering record is entirely in these details.

Accuracy is a specification, not an adjective

A survey-grade underwater photogrammetry system should be able to state its accuracy as a number, typically expressed as millimetres of error per metre of range, under stated conditions (visibility, distance, lighting, platform stability). "High-resolution" and "detailed" are not accuracy claims. If a vendor cannot give you a number and the conditions it applies under, ask why — and ask whether that figure has been independently certified rather than self-reported. DroneQ Robotics' Voyis Discovery Stereo system, for example, quotes 0.1mm/m accuracy under a Bureau Veritas certified inspection rating, which is the kind of independently verifiable claim worth asking every vendor for.

Repeatability matters more than any single survey

A single accurate 3D model is useful. A repeatable capture process — one that produces directly comparable models on the next visit, with the same referencing and the same measurement definitions — is what actually enables trending, and trending is what most integrity management frameworks are built around. Ask how a vendor aligns repeat surveys to the same coordinate reference, and how they quantify change between two visits, before you ask about resolution.

Georeferencing and traceability

A model that isn't accurately georeferenced to the asset's own coordinate system is a nice picture, not a survey deliverable. For anything feeding into class society or insurer documentation, you also want a clear chain of custody: raw data, processing method, and QA step, so the deliverable can be defended if it's ever challenged. A processing report that documents reprojection error, alignment accuracy, camera count and georeferencing residuals — not just the final pretty model — is a reasonable minimum to ask for.

The honest caveat: 3D imaging is not always the right tool

- Very low visibility water can defeat any camera-based system; acoustic methods (multibeam, structured-light-with-sonar hybrids) may be more appropriate, at lower resolution.
- 3D visual data tells you about geometry and surface condition — it does not tell you wall thickness, cathodic protection potential, or whether a member is flooded. Those require NDT sensors (ultrasonic testing, CP probes, flooded member detection), which can be integrated onto the same ROV mobilisation but are a genuinely different measurement.
- A single 3D survey without a repeat program behind it gives you a very good baseline and not much else. The value compounds on the second and third visit, not the first.

4. A PRACTICAL PRIMER: HOW UNDERWATER 3D PHOTOGRAMMETRY WORKS

You do not need to become a photogrammetrist to manage a 3D inspection program, but a working understanding of the pipeline helps you ask better questions of any vendor.

The capture step

A stereo or multi-camera rig, mounted on a ROV (or occasionally a diver-held frame), captures overlapping images or continuous video of the target as the vehicle moves along a planned survey path. Lighting, camera synchronisation, and vehicle stability during capture all directly affect the accuracy of everything downstream — this is why platform choice matters: a stable, well-instrumented work-class ROV will out-perform a lightweight inspection-class vehicle in current or poor visibility, even with a comparable camera.

A note on global shutter vs. rolling shutter

Conventional underwater cameras often use rolling shutter sensors, which expose an image line by line rather than all at once. When the ROV pitches, rolls, or moves through current, this can introduce motion-induced geometric distortion that quietly degrades photogrammetric accuracy without being obvious in the raw footage. Global shutter sensors, which expose every pixel simultaneously, eliminate this failure mode — a meaningful practical advantage in dynamic underwater conditions, not just a spec-sheet detail.

The processing step

Overlapping images are matched and triangulated to produce a dense point cloud — a set of individually measured 3D points, typically delivered as LAS/LAZ, PLY, XYZ or E57 files. That point cloud can be converted into a 3D mesh (a continuous triangulated surface, exported as OBJ, FBX, STL, PLY or GLB/glTF) or an orthomosaic (a flat, scale-accurate composite image, typically a GeoTIFF, that you can measure directly from). A textured 3D model — a mesh with photographic texture mapped onto it — is often the most useful format for visual inspection review, since it looks like the asset rather than an abstraction of it.

The measurement step

This is where a lot of programs stall. A model on its own doesn't answer an engineer's question; someone or something needs to extract measurements from it — wall dimensions, scour depth, marine growth volume, defect size — ideally without shipping the raw data back and forth by file transfer every time a stakeholder wants a number. This is the argument for a hosted, cloud-based delivery model: a place where measurement tools can be applied directly to the model by the client's own engineers, on demand, with outputs like distances, areas and volumes exported as CSV or JSON rather than requested by email.

The review step

Manually reviewing every frame of every survey for anomalies does not scale, and most of the data has nothing wrong with it. AI-assisted anomaly detection — trained to flag deviations from expected geometry or surface condition, and to classify them by type and severity — exists specifically to cut a large capture down to a short, prioritised list a human actually needs to look at. Ask any vendor proposing this how false positive and false negative rates are managed, and whether a human reviews flagged items before they reach you; a good AI-assisted workflow augments an engineer's judgement, it doesn't replace it.

Comparing capture methods

Method	Strengths	Limitations
Stereo/multi-view photogrammetry	Sub-millimetre accuracy achievable; rich visual + dimensional data; works with standard ROV mobilisation	Needs reasonable visibility; accuracy degrades with range and turbidity
Laser/structured-light scanning	Very high point density at close range; less affected by some lighting conditions	Shorter effective range; slower coverage of large areas; typically higher equipment cost
Multibeam / acoustic imaging sonar	Works in zero-visibility water; long range; good for large-area context	Much lower resolution than optical methods; not a substitute for detailed defect-level measurement
2D video + diver/pilot narrative (status quo)	Fast, familiar, low up-front cost, works everywhere	No persistent measurement; not repeatable/comparable; reviewer-dependent; poor archive value

In practice, the strongest programs we've seen use optical photogrammetry as the primary detailed-inspection method, with acoustic methods reserved for low-visibility conditions or large-area context surveys, and NDT sensors layered in wherever the specification calls for quantitative wall-thickness, CP or flooded-member data.

A worked example: two sensor classes, compared honestly

It's worth illustrating this with real, published figures rather than marketing language, using two systems at opposite ends of the offshore/nearshore spectrum. A larger-sensor offshore stereo system (of the kind fitted to a work-class ROV) may offer a theoretical maximum point density around 340 points/cm² in high-accuracy mode, dropping to around 44 points/cm² in a faster real-time mode suitable for coverage monitoring but not certifiable metrology.

A more compact, portable stereo scanner (of the kind fitted to a lighter inspection-class ROV) may offer around 90 points/cm² raw, with a wider field of view that covers a larger footprint per frame — useful in confined or complex geometry. Both approaches can converge on similar final accuracy once multi-frame photogrammetric processing is applied, but their raw single-frame density and practical field-of-view trade-offs are genuinely different, and worth asking any vendor to explain for your specific use case rather than taking a single headline accuracy number at face value.

5. BEYOND THE PRETTY PICTURE: MAKING 3D DATA USEFUL

The single biggest predictor of whether a 3D imaging program delivers value is not the accuracy of the camera. It's whether anyone at the client organisation actually uses the model after it's delivered. That depends on a handful of practical decisions made before the first survey, not after.

Decide where the data will live

A model sitting in a folder on someone's laptop, in a proprietary format only the vendor's software can open, will get used once and then forgotten. A model hosted somewhere the client's own engineers can open it, measure from it, and pull a report from it — without waiting on the vendor — gets used repeatedly. Modern cloud delivery models go further than file hand-off: point clouds and meshes can be published as 3D Tiles (a streaming format originally developed for large-scale geospatial data) or as browser-based WebGL viewers, so a stakeholder can review a model with nothing more than a web browser — no GIS software, no plugin, and no round trip offshore. If you're evaluating vendors, ask specifically where the data lives after delivery, in what formats, and who can access it.

Decide what format you actually need

A mature delivery set typically spans several product types, each with established file formats:

- Dense point cloud — LAS/LAZ (the most common format for subsea engineering workflows), PLY (includes per-point colour), XYZ/CSV, or E57.
- 3D mesh — OBJ (with texture files), FBX (for engineering/simulation tools), PLY, STL (geometry only), or GLB/glTF (lightweight, for web viewing).
- Textured 3D model — a mesh with photographic texture mapped on, typically OBJ or FBX with texture files, or a single-file GLB for browser viewing without local software.
- Orthomosaic — a georeferenced, distortion-free top-down image, usually delivered as GeoTIFF for engineering/GIS use, with JPEG/PNG for lightweight visualisation.
- Digital elevation model (DEM/DSM) — a raster elevation grid, usually GeoTIFF, used for scour, seabed change and volumetric trending.
- Supporting data — camera position/navigation logs, control-point/georeferencing residuals, and measurement/annotation exports, typically as CSV, JSON or XML, plus a processing/QA report (PDF or HTML) documenting reprojection error, alignment accuracy and georeferencing residuals.

Most integrity teams end up needing more than one of these formats for different audiences within the same organisation — a GIS engineer wants GeoTIFF, a reliability engineer wants the point cloud, and a manager wants something that opens in a browser without an explanation. Confirm upfront that your vendor can deliver in the formats you'll actually need, rather than a single proprietary export.

Plan for the second survey before you commission the first

A single 3D survey is a baseline. Its value roughly doubles the moment there is a second, comparable survey to measure change against, and compounds from there. If your specification doesn't already define a repeat cycle (annual, biennial, condition-triggered), decide on one before the first mobilisation — retrofitting a repeat program onto a one-off survey, with a different vendor or a different capture method, often means the first survey can't be used as a clean baseline.

Where NDT fits

For integrity-critical structures under an API- or NORSOK-style inspection specification, 3D visual data is usually necessary but not sufficient. Ultrasonic Testing (UT) for wall-thickness and corrosion mapping, Cathodic Protection (CP) measurement for protection-system verification, and Flooded Member Detection (FMD) for water ingress in tubular members are quantitative NDT techniques that answer different questions than a camera can. Where the same ROV mobilisation can carry both the 3D imaging payload and NDT probes, you save a second mobilisation; where it can't, plan the two scopes together rather than as an afterthought.

Budget for training, not just data capture

A hosted platform with measurement tools and AI-assisted anomaly review is only as useful as your team's ability to use it. Structured onboarding — a few hours introducing the platform, the measurement tools, and how to interpret AI-flagged anomalies — is a small cost relative to the survey itself, and is the difference between a model that gets used and one that gets filed away. If a vendor doesn't offer this as standard, budget time internally to build it yourself.

6. TEN PRACTICAL QUESTIONS TO ASK BEFORE YOU COMMISSION A SUBSEA 3D SURVEY

These are the questions that, in our experience, separate a programme that delivers ongoing value from one that produces an expensive one-off report.

1. **What accuracy can you commit to, in millimetres per metre, and under what conditions?** Get the number, the conditions, and — ideally — an independent certification (a class society or similar third party), not just a sales conversation.
2. **How do you align repeat surveys to a common reference, so I can measure change over time?** If the answer is "we'll figure it out next time," that's a red flag.
3. **Where does the data live after delivery, and in what formats?** Point cloud, mesh, orthomosaic, elevation model — confirm you get open, standard formats, not a single proprietary export.
4. **Can my own engineers take measurements from the model without coming back to you every time?** This determines whether the deliverable is a live tool or a static report.
5. **If you use AI-assisted anomaly detection, how are findings verified, and by whom, before they reach me?** You want augmentation of human review, not a black box.
6. **Can NDT (UT, CP, FMD) be captured in the same mobilisation, or does that require a separate contractor and a separate trip?** This affects both cost and schedule materially.
7. **What training or onboarding is included so my team can actually use the deliverable?** Ask what happens after the first login, not just the first delivery.
8. **What's your platform's track record in the specific conditions I operate in — visibility, depth, current?** A vendor with only clear-water reference projects may not perform the same way in your conditions.
9. **Is your inspection process independently compliant with recognised offshore standards (e.g. IMCA), and is your accuracy claim independently certified?** Ask to see the certificate or manual reference, not just a description.
10. **What does a repeat-survey program cost and look like over three to five years, not just the first mobilisation?** A vendor who can answer this clearly is thinking about your asset's lifecycle, not just their next invoice.

7. COMMON PITFALLS WHEN ADOPTING 3D DIGITAL TWINS

1. Treating the first survey as the whole project

The pilot survey is often the easy part. The harder, more valuable part — repeat surveys, trending, integration into your integrity management workflow — is where most of the return sits, and where most programmes quietly stop after the initial deliverable lands.

2. Over-trusting accuracy claims that don't specify conditions

An accuracy figure quoted for a calm test-tank pilot may not hold in a 2-knot current at 40 metres depth in 3-metre visibility. Ask for accuracy claims tied to conditions similar to your actual operating environment, and treat marketing accuracy numbers with the same scrutiny you'd apply to any other engineering specification.

3. Letting data formats lock you into a single vendor

If your data only exists in a vendor's proprietary viewer, switching providers later means starting your visual record over. Insist on open, standard formats (LAS/LAZ or E57 point clouds, OBJ/FBX/GLB meshes, GeoTIFF orthomosaics) as part of every deliverable, even if you also receive a proprietary viewer as a convenience layer.

4. Skipping training and then blaming the technology

A recurring pattern: a capable 3D model is delivered, no one on the client side was trained to use the hosting platform's measurement tools, the report gets skimmed once, and the following year's budget conversation concludes that "3D didn't really help." The technology wasn't the problem.

5. Mixing capture methods between surveys without accounting for it

Switching from optical photogrammetry to a different sensor type between repeat surveys — or changing vendors without aligning referencing — can make direct comparison unreliable exactly when you need it most, at the point where you're trying to detect change.

8. A PHASED ADOPTION PATH FOR ASSET INTEGRITY TEAMS

Teams that get sustained value from subsea 3D imaging tend to follow a broadly similar path, whether or not they plan it explicitly.

Phase 1 — Pilot

A single survey on a representative, ideally moderately accessible asset, explicitly scoped as a pilot: agree the accuracy target, the deliverable formats, and how success will be judged before the vessel mobilises. Use this phase to test the vendor's process and your own team's ability to use the output, not just to get a pretty model.

Phase 2 — Baseline

Extend the same capture standard across the asset population you intend to manage this way (a wind farm's monopiles, a platform's jacket members, a port's berths). This is the point at which you're building a comparable reference set, not a one-off.

Phase 3 — Repeat programme

Commit to a defined survey cycle (annual, biennial or condition-triggered) using the same capture and referencing standard, so that every subsequent survey adds to a trended, comparable record rather than starting again.

Phase 4 — Integration

Connect the hosted 3D data and its measurement/anomaly outputs into your broader integrity management workflow — CMMS, APM or digital twin systems — via exported CSV/JSON data or an API, so that 3D survey findings inform maintenance planning and risk assessment directly, rather than sitting in a separate silo from the rest of your asset data.

> *Most of the value in subsea 3D imaging is realised in Phases 3 and 4, not Phase 1 — plan your budget and your vendor relationship with that in mind from the outset.*

9. WHERE DRONEQ ROBOTICS FITS

DroneQ Robotics delivers subsea 3D imaging as part of Advanced Unmanned Robotics Services (AURS): worldwide ultrahigh-resolution inspection, survey, 3D Imaging and NDT, delivered underwater, on the water and from the air, for clients across ports, Offshore Oil & Gas and Offshore Wind. Subsea operations run under the IMCA-compliant DroneQ Robotics Subsea Inspection Operations Manual, which is routinely updated to reflect current techniques, technologies and operational insight.

ROV Platforms. Two ROV platforms especially adapted for 3D Digital Imaging sit behind that service.

The Argus Rover (OCROV Class II) carries the Voyis Discovery Stereo imaging system, rated to 4.000m and capable of capturing hairline coating defects and early-stage corrosion at 0.2mm/m accuracy — a independently certified rating, not a self-reported one. Its VSLAM-based real-time 3D modelling and Deep Vision Optics processing remove the visual effects of water to produce true-colour, sharp imagery with optimised depth of field across both close-up and distant features.

The compact X-SV (OCROV Class I) carries the UVision UScanner, built around a global shutter stereo camera pair that eliminates the motion-induced distortion rolling-shutter cameras introduce, combined with inertial and depth sensing for 1 millimetre-scale reconstruction. Its OctoDrive eight-thruster vector layout gives full 360° omnidirectional movement for the difficult-to-reach spots larger platforms can't access, and its onboard processing generates a live point cloud during the scan itself, so gaps in coverage are caught before the ROV leaves the water rather than on the next mobilisation.

DroneQ 3D Cloud. Survey output is delivered either as files (dense point clouds in LAS/LAZ, PLY, XYZ or E57; 3D meshes in OBJ, FBX, STL, PLY or GLB) or through DroneQ Robotics' cloud delivery model DroneQ 3D Cloud, where point clouds, textured 3D mesh, orthomosaics, dynamic elevation and Gaussian Splats models are hosted, processed and made available as downloadable files or as streaming 3D Tiles and browser-based WebGL viewers — no GIS software required to review a model. Clients can take their own dimensional and volumetric measurements, Construction Level measurements, run Depth Scan and Elevation Map analysis, and review AI-driven Anomaly Annotation — findings automatically classified by Anomaly Type and Severity rating — directly in DroneQ 3D Cloud. Structured User Training is provided at onboarding so that a client's own engineers can use the platform independently, rather than depending on DroneQ Robotics to interpret every deliverable.

None of this changes the guidance in this document. The questions in chapter 6 apply to DroneQ Robotics exactly as they apply to any other provider you evaluate, and we would rather earn a repeat-survey relationship by answering them well than by being the loudest voice in the room.

GLOSSARY

- Dense point cloud: a set of individually measured 3D points captured from overlapping survey imagery, typically delivered as LAS/LAZ, PLY, XYZ or E57.
- 3D mesh: a continuous triangulated surface generated by connecting points in a point cloud, typically delivered as OBJ, FBX, STL, PLY or GLB/glTF — useful for visualisation and volumetric measurement.
- Textured 3D model: a mesh with photographic texture mapped onto its geometry, giving an inspection-ready model that visually resembles the surveyed asset.
- Orthomosaic: a geometrically corrected, scale-accurate composite image of a surveyed area (typically GeoTIFF), allowing direct measurement from the image itself.
- Digital elevation model (DEM/DSM): a raster model of surface elevation over an area, used to track scour, seabed change or volumetric trends over time.
- 3D Tiles / WebGL streaming: cloud-native formats that let a 3D model be viewed and navigated directly in a web browser, without local GIS or CAD software.
- VSLAM: visual simultaneous localisation and mapping — a technique that builds a 3D model and tracks camera position in real time from visual input alone.
- Global shutter: a camera sensor design that exposes every pixel at the same instant, eliminating the motion-induced distortion that rolling-shutter sensors can introduce underwater.
- UT (Ultrasonic Testing): a non-destructive technique measuring material wall thickness to identify corrosion or erosion loss.
- CP (Cathodic Protection) measurement: verifies a structure's corrosion-protection system is within its design protection-potential range.
- FMD (Flooded Member Detection): identifies water ingress into tubular structural members, an early indicator of structural integrity loss.
- IMCA: the International Marine Contractors Association, whose guidance is a de facto industry standard for offshore ROV and diving operations.
- Survey-grade: capture accuracy and repeatability sufficient to support engineering measurement and change-detection, typically specified in millimetres of error per metre of range under stated conditions, ideally independently certified.

ABOUT THIS GUIDE

This guide was produced by DroneQ Robotics as a practical reference for asset integrity, inspection planning and ROV/marine coordination teams evaluating subsea 3D digital imaging. It reflects operational experience across Ports, Offshore Wind and Offshore Oil & Gas inspection programs delivered under DroneQ Robotics' Advanced Unmanned Robotics Services (AURS). Feedback, questions or requests for a working session on your own inspection specification are welcome via the contact details below.

DroneQ Robotics B.V.
Slotenmakerstraat 3
6101SK Echt (L), The Netherlands
info@droneq.nl
www.droneq.nl