WHITE PAPER
HLX-WPR-2026-014 · Rev B
August 2026

WHITE PAPER

How to Specify a Wall-Loss Screening and Continuous Monitoring Program

Criteria for buying integrity data your engineers can act on — what to demand from any vendor, what every method honestly cannot do, and when continuous monitoring earns its place.

SHOWN: AGING UTILITY-TUNNEL PIPING — THE ASSETS THIS PAPER IS ABOUT

About this paper

Most wall-loss programs are specified around an instrument. This paper argues they should be specified around the data: how often it arrives, whether it lands somewhere your team can find again, and whether its limits are stated plainly enough to engineer against.

It is written for the corrosion engineer, reliability engineer or inspection coordinator drafting an RFQ for screening or monitoring of piping, vessels and structural assets. It compares the common measurement methods — including where each one, HELIX's included, falls short — and closes with a fifteen-line specification checklist you can paste into a procurement document.

IN THIS PAPER

1. Why periodic inspection under-samples degradation
2. Screening methods compared — what each measures, and what it cannot
3. What to demand from spatial data — registration, not spreadsheet rows
4. From screening to continuous monitoring — when a permanent sensor earns its place
5. The specification checklist — fifteen criteria for an RFQ
6. Reference frameworks

Sections 1–4 build the argument; Section 5 compresses it into requirements you can paste into a procurement document; Section 6 anchors the terminology in the reference frameworks an RFQ can cite.

Prepared by Mateo Zane · HELIX Sensors Inc. · August 2026

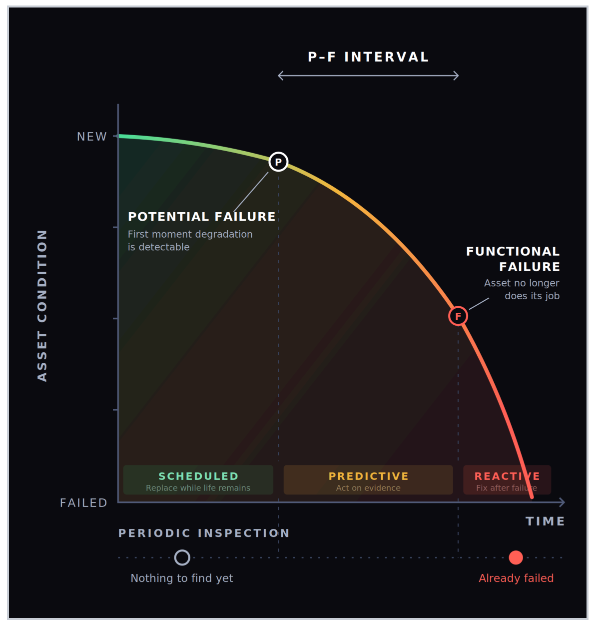

The P-F interval: degradation becomes detectable at P and ends function at F. A periodic program only catches it if a visit lands inside that window.

Executive summary

Wall loss is a continuous process. Almost every program that manages it is not: it samples the asset at intervals set by budget, access and the calendar, then holds the result as the truth until the next visit. Between visits, the wear rate is an assumption — and when the assumption is wrong (a process change, a failed coating, an erosion mechanism nobody predicted), the program finds out at the next inspection or at the failure, whichever comes first.

Specifying a better program is not mainly a question of instrument selection. The differences that decide whether the data is usable five years from now are structural, and they belong in the RFQ:

- **Honesty about the measurement.** A vendor who states the footprint, the detection floor and the failure modes of their own measurement is giving you something you can engineer against. One who does not is selling you a colour.
- **The cost of access, counted properly.** For insulated, elevated, buried or congested assets, access — scaffolding, insulation removal, surface preparation, permits, shutdowns — dominates the cost of every reading. Methods that work through coatings, in service, change the economics more than any difference in gauge accuracy.
- **Location repeatability.** A thickness trend is only as good as the certainty that scan two measured the same steel as scan one. Demand a mechanism, not a promise.
- **Registration to the asset, not to a spreadsheet.** Data keyed to row numbers and drawing references decays with staff turnover. Data registered to a 3D model of the physical asset stays findable, comparable and auditable.
- **Continuity where it is justified — and only there.** Permanently installed monitoring earns its place on a shortlist of locations where access is expensive, consequence is high, or the wear rate is faster or less explained than the inspection interval can track. It is the wrong tool for random pitting and for one-off surveys.

The stakes are not abstract. ABB's Value of Reliability survey puts unplanned downtime near C$242,000 per hour for a typical Canadian industrial business. One operator we work with reports $3.5M in downtime cost on a single two-week inspection cycle, repeated every three years — spent largely on gaining access to take readings whose result, most years, is "no significant change."

This paper sets out the argument in six short sections and ends with a checklist. Where HELIX's own methods have limits, they are stated in the same terms we suggest you demand from any vendor — including us.

> *A specification that asks every vendor to state what their measurement averages, what it cannot resolve, and how a reading will be found again in five years will do more for your program than any single instrument choice.*

1 Why periodic inspection under-samples degradation

The standard model of condition monitoring is the P-F curve: an asset degrades, the degradation becomes detectable at some point P, and the asset stops doing its job at point F. The window between them — the P-F interval — is the entire opportunity a program has to act on evidence rather than on failure. For wall loss, that window can be years or months: erosion at a slurry bend, corrosion under a failed coating and flow-accelerated corrosion after a process change all move the F point without notice.

A periodic program samples this continuous process on two axes at once, and under-samples on both:

- **In time.** If inspections land every one to five years, any mechanism whose P-F interval is shorter than the gap between visits is invisible until it is late. Two readings years apart give one rate estimate — with no confidence bounds, and no way to tell steady loss from accelerating loss or to say when the rate changed.
- **In space.** A spot reading measures the point under the probe, and grids cover a small fraction of the surface, chosen where damage is expected. Damage that concentrates somewhere else — at a support contact, under a drip, at the extrados of one bend out of forty — is sampled only if someone thought to put a point there.

A third failure mode is administrative rather than physical: the reading exists, but the program cannot confidently relocate it. If the follow-up campaign cannot re-occupy the same position, the "trend" between two campaigns includes an unknown offset from measuring different steel — quietly absorbed as scatter.

The honest counter-argument

The textbook fix is to inspect more often, and where access is cheap that is often the right answer — a walkable, uninsulated line at grade does not need permanent instrumentation to be well managed. The economics turn when access dominates: insulation removal and reinstatement, scaffolding, rope access, confined-space or excavation permits, and production interruption can each cost more than the measurement itself. At that point, each increment of confidence is bought at close to full campaign cost.

WHAT THE INTERVAL COSTS *(third-party and operator-reported figures — sources below)*

C$242,000	$3.5M	3 years
PER HOUR — UNPLANNED DOWNTIME, TYPICAL CANADIAN INDUSTRIAL BUSINESS	DOWNTIME ON ONE TWO-WEEK INSPECTION CYCLE, REPORTED BY ONE OPERATOR	BETWEEN THOSE LOOKS AT THE ASSET'S CONDITION

Sources: ABB Value of Reliability survey (downtime cost); operator figure as reported to HELIX by one operator we work with — a single site's experience, not an industry average. Your own access and downtime costs are the numbers that should drive the specification.

The rest of this paper turns that framing into procurement criteria: first for the screening methods that establish where an asset stands today, then for the data infrastructure those results should land in, and finally for where continuous monitoring is worth installing.

2 Screening methods compared — what each measures, and what it cannot

No single method covers wall loss. A defensible program pairs a wide, fast screening method — to find where attention belongs — with precise follow-up where it matters. The comparison below states each method's honest limits in the terms we suggest you require from any bidder. Two of these methods are HELIX's own; their limits are listed just as bluntly.

Method	What it actually measures	Access needed	Where it stops
Visual inspection	External surface condition: coating breakdown, blistering, staining, visible corrosion product, mechanical damage	Line of sight to the surface	Qualitative only. Sees nothing through insulation or cladding, and nothing about remaining wall — a sound-looking coating can hide advanced loss, and surface rust can overstate it
UT spot readings	Wall thickness at the point under the probe — a direct, absolute measurement, with a mature certified workforce behind it	Bare, prepared metal and couplant at every point; insulation windows or coating removal where present	Samples only the points measured. Trend quality depends on re-occupying each point exactly; grids cover a small fraction of the surface, and each point carries its access cost on every repeat visit
PEC screening (HELIX PEC)	Decay of a pulsed eddy-current response — a relative apparent-wall value averaged over a ~90 mm footprint, trended and compared within consistent material and geometry	None of the above: reads through coatings and insulation, no surface preparation, no couplant, asset in service	Not an absolute gauge without a calibration reference. Footprint-averaged — will not resolve isolated pitting or tight cracking. Responds to deposits, scale and ferromagnetic corrosion product as if they were wall. Sensitivity falls as liftoff grows
ECA screening (HELIX ECA)	Surface-breaking and near-surface discontinuities — weld toe, cap and heat-affected-zone indications — from a 32-element array scanned through paint	A path the encoder can be run along; coating stays on	Surface and near-surface only. Reports indications; it does not size defects. Heavy scale or an irregular surface profile degrades the signal
3D surface capture (LiDAR + photogrammetry)	Three-dimensional external geometry of the asset, resolving individual pitting on exposed surfaces	Line of sight	Surface only — measures nothing through the wall. Its role here is as the spatial skeleton the other measurements are registered onto (Section 3)

What to put in the RFQ. Require each bidder to state, in writing: the measurement footprint and what falls below it; whether values are absolute or relative, and what reference makes them absolute; the known confounders (deposits, scale, liftoff, temperature); and the detection floor *as evidenced in field data, not on a calibration block*. A method with honestly stated limits can be engineered around. An unstated limit is where programs get surprised.

Note what this table implies for sequencing: screening methods that read through coatings buy coverage cheaply and point precise methods at the right places. Spot UT remains the right tool for absolute verification at locations screening has flagged — the two are complements, not rivals.

3 What to demand from spatial data — registration, not spreadsheet rows

The default deliverable of an inspection campaign is a report PDF and a spreadsheet: one row per reading, keyed to a location ID and a drawing reference. The data is real, but its usefulness decays: location IDs drift from as-built reality, the engineer who knew what "Z14-R27" meant changes roles, and the next campaign lays out a different grid the old one can no longer be compared with point-for-point. Measurements paid for at full access cost function, within a few years, as history rather than as a baseline.

The alternative is to make the asset itself the index. If the campaign captures a 3D model of the physical asset — LiDAR plus photogrammetry — and every reading is registered to its position on that model, the data stays interrogable: any engineer can find a reading by looking where it physically lives, repeat campaigns land on the same geometry, and comparison becomes an overlay instead of a spreadsheet-matching exercise. This is what HELIX Horizon does, and we would make the same argument about any system that does it: registration is a property of the data, and it belongs in the specification.

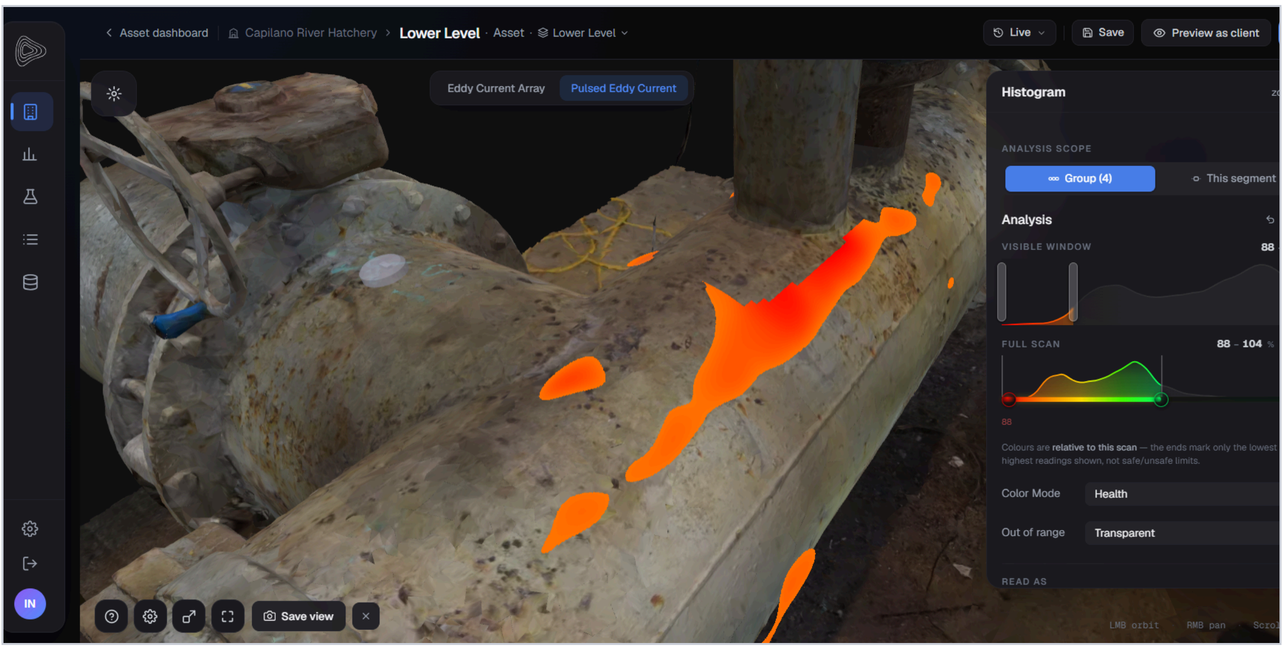

Screening registered to the asset in HELIX Horizon: PEC layer on the Capilano River Hatchery (DFO) site model, zone histogram open. Colours are relative to this scan — they mark the lowest and highest readings shown, not safe/unsafe limits.

Concretely, registration is what makes the following possible — and each is a line you can demand in an RFQ regardless of vendor:

- **Point-level recall:** hover any location and read its value; findings carried as a point-of-interest register pinned to the geometry, not as prose in a report appendix.
- **Zone-level statistics:** per-zone histograms and distributions, so "worst 5% of this section" is a query, not a manual sort.
- **Repeat-scan comparison:** scan history on a time slider over the same geometry, so change between campaigns is visible where it happened.
- **Honest colour handling:** explicit display modes for values outside the visible window, and a stated meaning for the colour scale on every view.

Scale of proof: on a four-day screening of ~100 m of 1970s piping at a Fisheries and Oceans Canada hatchery, 378,407 screening samples across 34 zones were registered onto the site model, with 46 point-of-interest entries — screened in place, without cutting, draining or dismantling (HELIX case study HLX-2026-001).

4 From screening to continuous monitoring — when a permanent sensor earns its place

Screening establishes where an asset stands and where its worst locations are. The remaining question is which of those locations justify watching continuously — and the honest answer is: a shortlist. Permanently installed monitoring — in HELIX's case, a patent-pending flexible sensor wrap that bonds to the asset and trends wall thickness with temperature compensation — is a targeted instrument, not a blanket. It should be specified where these conditions overlap:

Where continuous monitoring earns its place

- **Access is the dominant cost.** Every repeat reading requires scaffolding, insulation removal, permits, excavation or a shutdown — so each increment of confidence costs close to a full campaign.
- **Consequence is high.** Failure means lost production, environmental release or a safety exposure that makes "wait for the next campaign" an uncomfortable sentence to write down.
- **The wear is fast or unexplained.** Screening found loss whose rate or mechanism is not established — the exact case where two readings years apart cannot tell steady from accelerating.
- **Replacement is scheduled by calendar.** Components changed on fixed intervals, where a measured rate would let replacement be timed to condition instead.
- **A retirement limit is approaching** and the decision needs a rate with confidence bounds, not an extrapolation from sparse points.

Where it does not

- **Random, isolated pitting.** A fixed grid of footprint-averaged readings will not resolve it; pitting is resolved only where it is large enough relative to the sensing footprint to move the average.
- **Unpredictable damage locations.** If screening cannot say where degradation concentrates, a permanent sensor cannot be scoped to cover it.
- **One-off surveys and cheap-access assets.** Handheld screening or spot UT on a walkable line is more economical, and remains so.

The premise, stated plainly. A monitoring installation does not slow wall loss by a single micron. What it changes is what the operator knows — and therefore how much conservatism has to be carried in inspection frequency, replacement scheduling and operating decisions. If that conservatism is not costing you anything, monitoring will not pay for itself.

The two halves of the program feed each other: screening nominates the shortlist, the installed sensor watches it, and both datasets live on the same registered model — so the point a wrap is trending is the same physical point the screening campaign flagged, and the next screening pass checks the area around the wrap that fixed sensors do not cover.

What continuous data should deliver: a rate, a forecast, and its confidence

The output that justifies a monitoring installation is not a thickness number — it is a defensible rate of change and a forward statement with confidence attached. Specify the form of that statement. It should read like an engineering finding: *no crossing of the criterion line before a stated date at a stated confidence; loss rate bounded below a stated value*. And the method behind it should be disclosed: HELIX Horizon's forecast is statistical trending of the measured rate of change with its uncertainty propagated forward — not a learned model, and we state that because you should ask every vendor the same question.

SIMULATED SCENARIO — DEMONSTRATION DATA, NOT A CUSTOMER RESULT

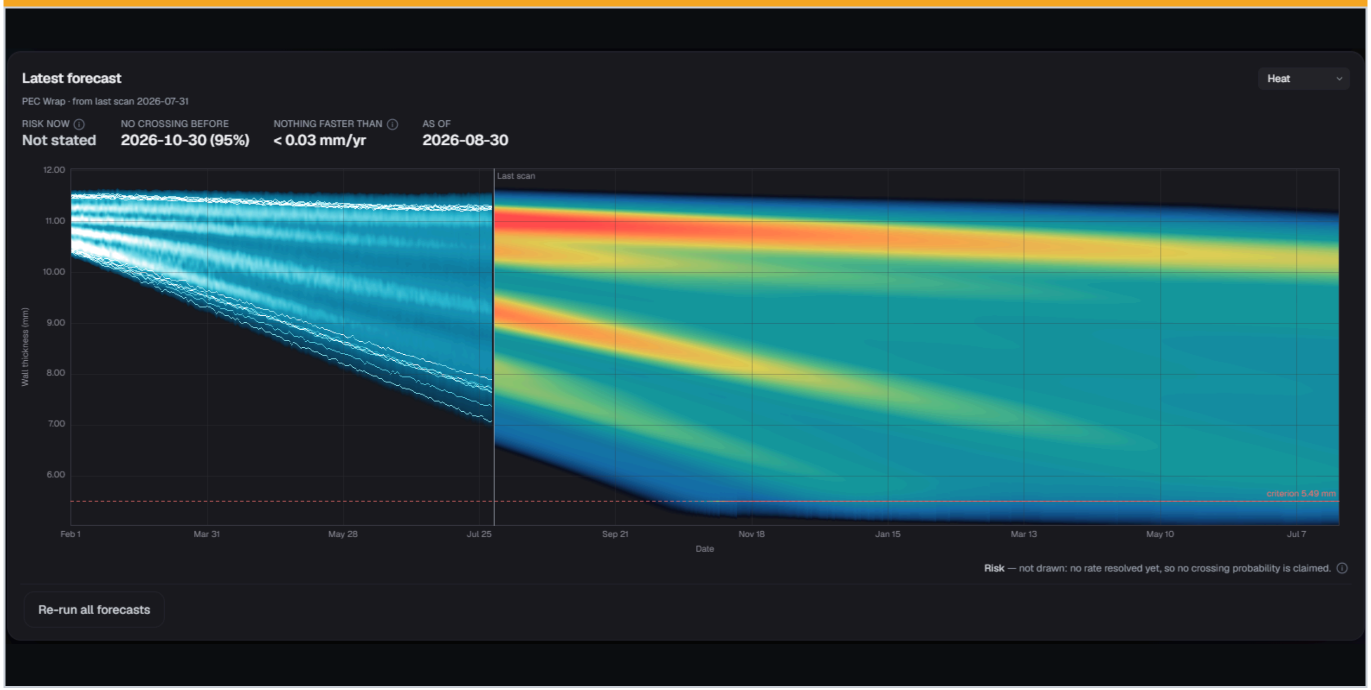

HELIX Horizon "Latest forecast" on a simulated erosion scenario: measured history to the last scan, then the probabilistic forecast fan — "No crossing before 2026-10-30 (95%)", loss rate <0.03 mm/yr against the criterion line. Simulated scenario data for demonstration.

Underneath a forecast like this sits measurement stability, which is the specification item that makes or breaks it: a drift of tens of microns that correlates with weather is indistinguishable from wall loss unless temperature is measured and compensated at the sensing point. Ask for the compensation method and the residual stability in writing. HELIX's published live wrap pilot record is the evidence we offer for ours:

LIVE WRAP PILOT — PUBLISHED FIELD RECORD *(HELIX "Array Live", helixsensors.com)*

12	3.000 mm	4,800	<1%	±25 µm
MEASUREMENT POINTS	NOMINAL WALL MONITORED	READINGS OVER 8 HOURS	WALL-LOSS DETECTION THRESHOLD	TEMPERATURE-COMPENSATED

A candid note on evidence: the market leaders in installed monitoring can point to installed bases a decade old, and if you are evaluating them, you should ask to see that field data. HELIX's installed base is young. What we publish instead is a live field record and pilots scored against criteria fixed in writing before commissioning — measured against the site's own inspection data, not our reporting. We think that is the right standard to hold any vendor to, and Section 5 writes it into the checklist.

5 The specification checklist — fifteen criteria for an RFQ

Each line below is phrased as a requirement you can paste into a procurement document. None of them name a technology; together they describe a program whose data survives contact with the next five years.

#	Require the bidder to…	Why it belongs in the spec
1	State the measurement physics, the footprint over which each reading averages, and what falls below that footprint (isolated pitting, tight cracking).	Every method averages something. Unstated footprints are where programs get surprised.
2	State whether reported values are absolute or relative, and exactly what reference is needed to make them absolute.	Relative data is legitimate and useful for trending — but only if everyone reading it knows that is what it is.
3	List known confounders in writing: deposits, scale, corrosion product, liftoff, temperature, surface profile.	A confounder you know about becomes a procedure; one you do not becomes a false trend.
4	Evidence the claimed detection floor with field data, not calibration-block or laboratory results.	Lab conditions flatter every instrument. Ask every vendor — including HELIX — for their field record.
5	Quantify the access burden per reading or per position: surface preparation, coating or insulation removal, couplant, scaffolding, permits, shutdown.	Access, not instrument time, dominates lifecycle cost on hard-to-reach assets. Make bidders price it visibly.
6	Measure with the asset in service, and state temperature and coating limits where that is not possible.	Data that arrives only at shutdowns inherits the shutdown interval — the under-sampling problem in Section 1.
7	Take measurements without breaking containment or penetrating the process boundary.	Intrusive points add leak paths, permit load and their own inspection burden for the life of the asset.
8	Guarantee location repeatability by mechanism — fixed sensing positions, registration to geometry, or documented re-occupation procedure.	A trend built on approximately re-found points contains an offset nobody can quantify.
9	Deliver every reading registered to its physical position on a 3D model of the actual asset, with exportable views.	Spreadsheet rows decay with staff turnover and re-gridded campaigns; geometry does not move.
10	State the meaning of every colour scale, and display out-of-window values explicitly rather than clipping them.	A heatmap where red silently implies "unsafe" — or worse, hides what is beyond the scale — invites unsupported decisions.
11	Report rate of change with confidence bounds, and state the forecasting method (statistical trending vs. learned model) behind any forward claim.	Decisions are made on the rate and its uncertainty, not on a thickness snapshot.
12	Provide threshold and rate-of-change alarms, configurable per zone, routed to the plant historian or SCADA.	Continuous data nobody is alerted to is a periodic program with extra steps.
13	Confirm the customer owns all measurement data, with open export (API or historian integration) at no additional licence.	Do not specify a data hostage. Ownership should survive any change of vendor.
14	Put certification claims in writing: which standards are referenced versus certified against, and what personnel certifications are actually held.	Claims vary in precision. HELIX, for the record: standards referenced as frameworks, no certifications claimed — see Section 6.
15	Offer a pilot with pass criteria fixed in writing before commissioning, scored against evidence the customer holds.	A vendor confident in their measurement will let your data judge it. Falsifiability is the cheapest due diligence available.

Rows 1–4 are the measurement-honesty criteria of Section 2; rows 5–7 the access economics of Sections 1 and 4; rows 8–10 the registration argument of Section 3; rows 11–13 the continuous-data outputs of Section 4; rows 14–15 the evidence standard this paper argues for throughout.

6 Reference frameworks

The standards below are the frameworks HELIX's methods are developed with reference to, and the ones most useful for anchoring method and terminology in an RFQ. They are cited here as reference frameworks: HELIX does not claim certification or compliance to them, and we suggest requiring the same referenced-versus-certified distinction, in writing, from every bidder (checklist row 14).

Framework	Scope	Use in a specification
ISO 20669:2017	Pulsed eddy current testing of ferromagnetic metallic components	Method and terminology reference for PEC screening requirements
ISO 15548-1/-2/-3	Characterization and verification of eddy current examination equipment	Equipment characterization language for any eddy-current bidder
ISO 20339:2017	Equipment for eddy current examination — array probe characteristics and verification	Array-probe verification reference for ECA screening rows
ASTM E3052-21	Examination of carbon steel welds using eddy current array	Practice reference for ECA weld-examination scope
ISO 17643:2015	Eddy current examination of welds by complex plane analysis	Companion weld-examination reference
ISO 9712 / CAN/CGSB-48.9712	Qualification and certification of NDT personnel	Where your program requires code-mandated inspection, require certified personnel. HELIX does not hold ISO 9712 certification: HELIX screening is screening, and does not certify to any code scheme

Screening is not fitness-for-service. Nothing in a screening or monitoring program — HELIX's or anyone's — replaces fitness-for-service assessment, which is a regulated engineering activity that stays with the operator's engineers. The role of the data described in this paper is to make that engineering better informed and to point certified inspection where it is needed, not to substitute for either.

Conclusion

Specify the data, not the instrument. A wall-loss program stands or falls on four properties that belong in the RFQ: measurements whose limits are stated plainly enough to engineer against; an access burden counted honestly across the program's life; readings that land on the asset's geometry so they stay findable and comparable; and — on a deliberately short list of locations — continuity, with rates and forecasts that carry their confidence with them. The fifteen-row checklist in Section 5 is this paper in executable form. It does not name a vendor, and a strong bid from anyone should survive it.

HELIX built its screening service, sensor wrap and Horizon software around these criteria, and we are content to be evaluated against every row — including the ones that ask hard questions of a young company. Send us the checklist filled in with your asset's realities, and we will respond line by line, with a written answer where our current honest response is a limitation.

About HELIX Sensors

HELIX Sensors Inc. is a British Columbia company building continuous health monitoring for critical assets. The HELIX platform combines non-intrusive NDT screening (pulsed eddy current and eddy current array, with LiDAR + photogrammetry capture), a patent-pending permanently installed sensor wrap for continuous wall-loss monitoring, and HELIX Horizon — software that registers every measurement to a 3D model of the physical asset and trends it over time. Screening is the entry point; the wrap and software carry the program between campaigns.

PUBLISHED FIELD RECORD *(live counters, helixsensors.com, August 2026)*

2,002 h	3,200+ ft²	360M+	378,407
WRAPPED SCAN TIME	SCANNED	DATA POINTS	SAMPLES — DFO CAPILANO SCREENING

Field proof includes a four-day screening of ~100 m of 1970s piping at Fisheries and Oceans Canada's Capilano River Hatchery — screened in place, without cutting, draining or dismantling (case study HLX-2026-001) — and a live wrap pilot whose field record is published at helixsensors.com. Five provisional patents are filed covering sensor hardware, deployment methods, signal processing and predictive modeling.

Related documents

Document	Contents
HLX-SPC-2026-000	Technologies application guide — what each method does, and where it applies
HLX-PDS-2026-006 / -007	Product data sheets — HELIX PEC and HELIX ECA screening scanners
HLX-PDS-2026-008	Product data sheet — HELIX Sensor Wrap continuous monitoring system
HLX-PDS-2026-009	Product data sheet — HELIX Horizon software
HLX-CAS-2026-010	Case study — NDT screening & 3D condition mapping, Capilano River Hatchery (DFO)

HELIX Sensors Inc.
1774 Billington Rd, Bowen Island, BC, Canada
info@helixsensors.com · 236-334-7871 · helixsensors.com

To put the Section 5 checklist to work on a specific asset — or to have HELIX answer it line by line — book a meeting at helixsensors.com or write to info@helixsensors.com.