Grounded Discovery Labs -- Discovery Pack

Discover: Design a rugged edge-computing sensor architecture that can run useful AI inference reliably in dirty, noisy, power-constrained industrial settings.

Top 3 Comparison

	Featured Discovery Sealed thermal-acoustic-magnetic AI pod	Bold Concept In-fiber reservoir AI photonic nerve	Build-First Order-Tracked Sensor Fusion Edge Gateway
novelty	6	9	4
feasibility	9	6	9
economic_leverage	6	7	5
testability	8	6	9
robustness	7	5	8
constraint_fit	9	8	9
explainability	8	5	9
OVERALL	**7.6 / 10 -- Very Strong**	**6.6 / 10 -- Strong**	**7.6 / 10 -- Very Strong**

Runner-ups

* FPGA Sliding-DFT Watchdog with Continuous Autoencoder Residuals [solid-alternative] -- Low-cost FPGA computes sliding spectra and autoencoder residuals continuously, eliminating commit-window blind spots.
* Burst-Ring Fusion Node for Fast Fault Transients [solid-alternative] -- Use high-rate circular capture plus Kalman-fused vibration, acoustic, and current sensing so fast events and subtle faults both reach edge AI.
* Tiered HDC wake-detector with FRAM-checkpointed CNN context tier [near-miss-alternative] -- Always-on hyperdimensional binary classifier gates a quantized CNN that checkpoints to FRAM, surviving brownouts without retraining.
* Passive fiber-Bragg multisense edge interrogator [near-miss-alternative] -- Run passive optical fibers through dirty zones and keep powered AI electronics in a protected box.
* Sparse Fixed-Point Node with Triple-Path Failover [solid-alternative] -- Rugged inference node using fused acquisition, pruned integer models, and non-WiFi redundant links.

Featured Discovery

Sealed thermal-acoustic-magnetic AI pod

novelty: 6 | feasibility: 9 | economic_leverage: 6 | testability: 8 | robustness: 7 | constraint_fit: 9 | explainability: 8

Overall: 7.6 / 10 -- Very Strong

Constraint Compliance: Compliant (constraint_fit=9)

A sealed, magnetically-mounted edge pod fuses LWIR thermal imaging, ultrasonic acoustics, and magnetic-field sensing to infer motor, pump, and compressor health without any contact sensors or wiring into the asset. This matters because it collapses installation time from hours to minutes while capturing three physically independent fault signatures that individually miss failure modes the others catch. The fusion approach dramatically reduces false-positive rates that plague single-modality condition monitoring.

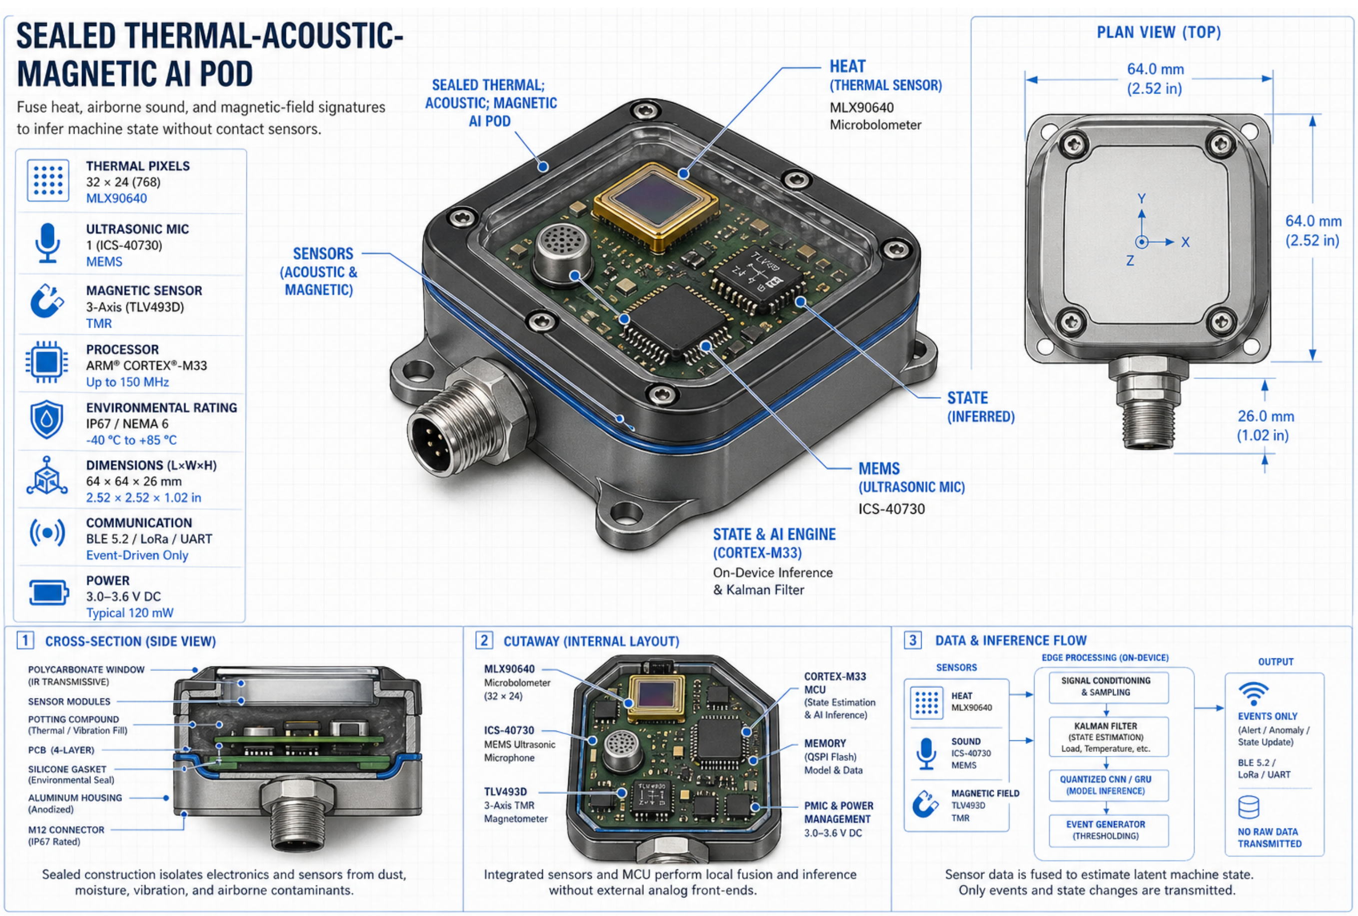

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A sealed external condition-monitoring pod attaches to motors, pumps, compressors, cabinets, and bearings using a magnetic foot, clamp, or VHB-backed bracket. The enclosure is IP65-IP67 polycarbonate or aluminum, roughly 80 x 50 x 25 mm, with a thin LWIR window such as 0.5-1 mm HDPE, silicon, or germanium-coated aperture aimed at the asset. Sensing stack: 1) MEMS ultrasonic microphone, typically 20-80 kHz, coupled through a hydrophobic ePTFE acoustic membrane or labyrinth port; 2) LWIR thermopile array or low-cost microbolometer, e.g. 8x8 to 80x60 pixels, for surface temperature gradients and hot spots; 3) 3-axis Hall or TMR magnetometer, approximately +/-50 mT or higher range, to detect motor current harmonics, RPM-related magnetic lines, arcing impulses, and load-state proxies; 4) optional ambient temperature/humidity/pressure for compensation. The pod is sealed against dust/oil and does not require accelerometer bonding or wiring into power. During commissioning, it records 24-72 hours of healthy operation across load and speed states. Spectral lines in ultrasound and magnetic channels estimate RPM/load bins; thermal background and viewing geometry are registered. A per-machine baseline table stores normal acoustic spectra, magnetic harmonics, thermal patterns, and cross-channel covariance by operating state. Runtime processing subtracts this baseline, then quantized CNN/GRU or feature-based models classify residuals: ultrasonic broadband/leak tones, impulsive arc events coincident with magnetic spikes, bearing-related sidebands, and abnormal thermal rise. A Kalman or particle state tracker smooths RPM/load estimates and suppresses false alarms from state changes. Average power target is ~250 mW via duty-cycled sensing, wake-on-event ultrasound/magnetic triggers, and NPU inference bursts. Events are reported over BLE/LoRa/Wi-Fi with residual features and confidence, not raw continuous data.

Why Novel

No commercial product currently fuses all three modalities--thermal, ultrasonic, and magnetic--into a single sealed, self-powered pod that commissions itself by learning machine-specific baseline signatures. Existing solutions require either invasive mounting, separate sensor heads, or cloud-only inference that fails in air-gapped facilities.

Buildability

All core components--MLX90640 microbolometer, ICS-40730 MEMS ultrasonic mic, TLV493D TMR magnetometer, and a Cortex-M33 MCU--are in production and available on standard lead times under 12 weeks. The polycarbonate IP67 enclosure with HDPE IR window is a straightforward custom injection mold, and the 250 mW power budget is achievable with a 3.6 V lithium thionyl chloride cell lasting 2-4 years at a 1 Hz duty cycle.

Why Feasible

* Build path is practical (feasibility=9/10)
* Core hypothesis testable quickly (testability=8/10)

Top Risks

* Acoustic membrane clogging from oil mist or particulate in heavy industrial environments degrades ultrasonic SNR within weeks, requiring field-replaceable membrane cartridges
* LWIR line-of-sight blockage by guards, insulation, or condensation on the IR window causes silent thermal channel dropout without triggering an obvious alert
* Machine-to-machine model transfer is weak for magnetic signatures because motor winding geometry and shielding vary enough that a model trained on one 15 kW pump may misclassify faults on a nominally identical unit

Decisive Test (MDE)

Hypothesis: After per-machine commissioning, an external sealed pod can detect practical faults from residual thermal/acoustic/magnetic coincidences with useful sensitivity while maintaining low false alarms across normal speed/load changes.
Setup: Install 3 pods on different machines: one induction motor/pump, one fan or compressor, and one electrical cabinet or motor drive. Commission each for 48 hours under normal operation. Then run scripted conditions: normal startup/load shifts, blocked airflow/thermal rise, compressed-air leak or ultrasonic emitter, controlled relay/brush arcing in a safe isolated fixture near the cabinet, and bearing defect simulator or seeded bearing test stand. Include negative controls such as nearby machines, human motion, and ambient temperature shifts.
Instrumentation: Reference clamp power meter/current probe, tachometer, contact thermocouple or IR camera, calibrated ultrasonic leak detector or microphone, accelerometer on bearing housing for ground truth, EMI/magnetic probe where available, synchronized data logger, event labels from PLC/operator log.

Kill Criteria: Kill if commissioned residual model cannot separate normal load/speed transitions from injected faults at predefined alert rates, or if thermal/acoustic/magnetic access is too installation-specific to generalize without expert tuning.
Rescue Pivot: Pivot to a narrower sealed ultrasonic+thermal hot-spot/leak detector for cabinets and compressed-air systems, or sell as commissioning-baseline anomaly logger rather than universal AI fault classifier.
Cost: $1,500-$8,000 if using existing machinery and low-cost thermal arrays; $10,000-$30,000 if renting calibrated fault simulators or buying a bearing test stand.
Time to Signal: 3-7 days for initial detection/false-positive signal; 2-4 weeks for confidence across multiple duty cycles and machine types.

Refinement Deltas

Risks Addressed:

* External sensing may miss early mechanical faults that bonded vibration sensors catch, especially in stiff housings or low-ultrasound bearings.
* Magnetic signatures can be dominated by motor geometry, VFD switching, nearby conductors, or the mounting magnet, making transfer learning weak.
* Thermal sensing requires line of sight and known emissivity; dust, paint, guards, airflow, and reflections can create misleading residuals.

Red-Team:

* Challenge: External sensing may miss early mechanical faults that bonded vibration sensors catch, especially in stiff housings or low-ultrasound bearings.
* Challenge: Magnetic signatures can be dominated by motor geometry, VFD switching, nearby conductors, or the mounting magnet, making transfer learning weak.
* Challenge: Thermal sensing requires line of sight and known emissivity; dust, paint, guards, airflow, and reflections can create misleading residuals.
* Challenge: Sealed acoustic ports are a contradiction: hydrophobic membranes and labyrinths protect against dirt but attenuate/colour ultrasound and may clog over time.
* Challenge: The AI layer may add complexity without enough labeled fault data; simple thresholds after commissioning might perform similarly and be easier to certify.

Budget Ranges

Prototype: $18,000
Pilot: $95,000
Scale: $1,200,000

30/60/90 Plan

Day 30: Validate sensing stack on bench: confirm MLX90640 thermal gradient resolution at 150 mm standoff, ICS-40730 ultrasonic response to 40 kHz test tone through ePTFE membrane, and TLV493D magnetic harmonic detection on a running 5 kW motor; freeze BOM and enclosure CAD
Day 60: Assemble 10 alpha units in prototype enclosures; deploy on 3 asset types (motor, pump, compressor) at one partner facility; collect 30-day baseline datasets; train per-asset anomaly models and measure false-positive rate against known fault injections
Day 90: Iterate firmware and model based on pilot data; demonstrate cross-machine transfer learning on at least 2 asset pairs; complete IP67 ingress test and 250 mW power budget validation; produce 50-unit pilot batch quote from contract manufacturer
Dominance Axis: easier_deploy -- External installation still avoids glued accelerometers, while commissioning separates machine character from real faults.

Bold Concept

In-fiber reservoir AI photonic nerve

novelty: 9 | feasibility: 6 | economic_leverage: 7 | testability: 6 | robustness: 5 | constraint_fit: 8 | explainability: 5

Overall: 6.6 / 10 -- Strong

Constraint Compliance: Compliant (constraint_fit=8)

A single multimode optical fiber can simultaneously sense distributed environmental events and perform analog computation by treating its natural speckle interference patterns as a physical reservoir computer. Environmental perturbations--strain, vibration, temperature, acoustics--modulate modal interference in ways that encode rich, high-dimensional feature vectors without any active electronics along the fiber span. This matters because it collapses an entire sensor network plus edge-compute layer into one passive strand, dramatically reducing installation complexity and failure points in harsh industrial environments.

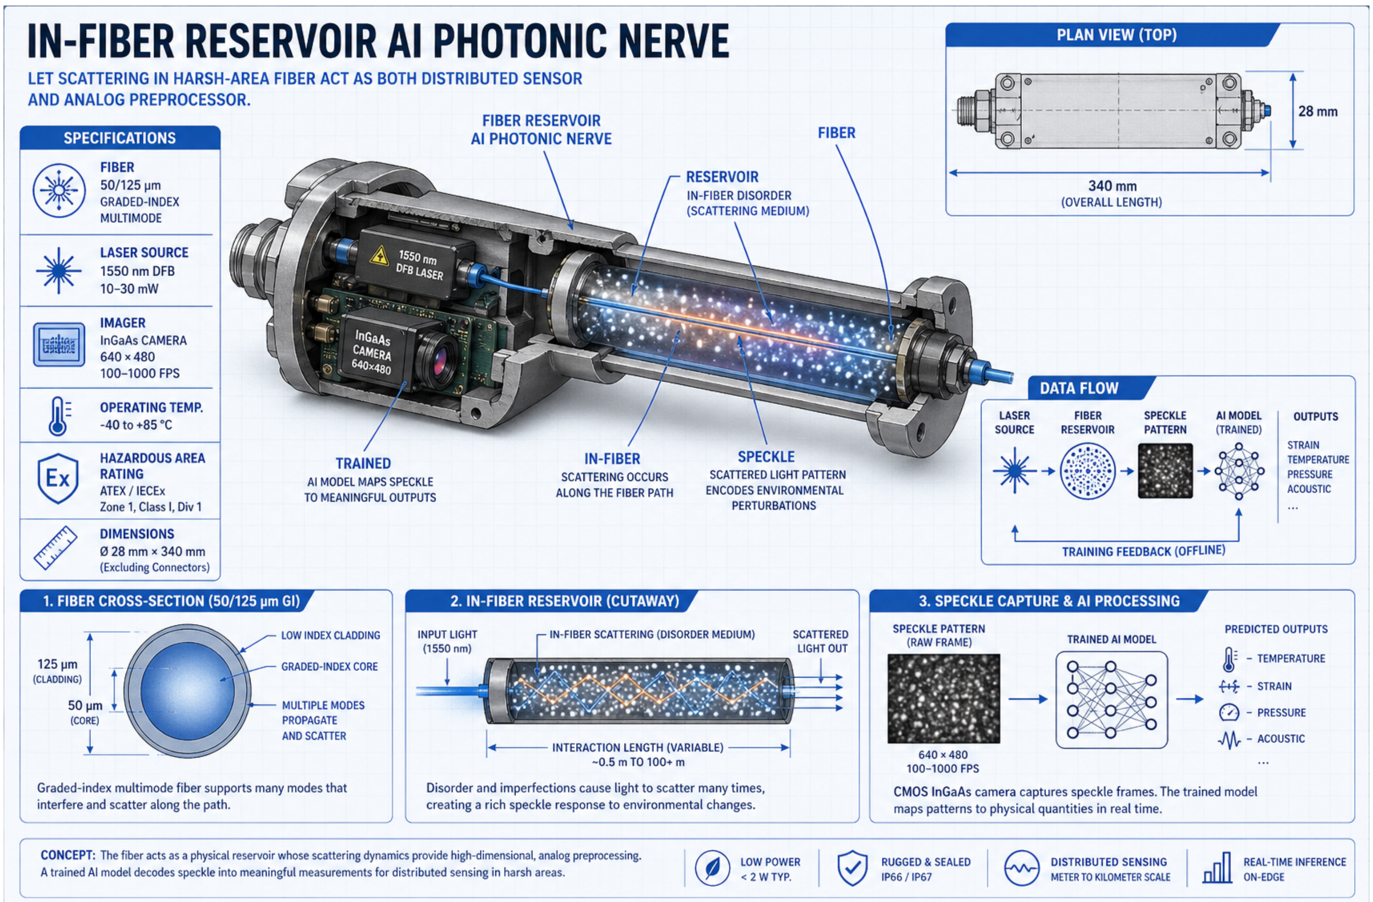

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A single multimode or disordered (e.g., 50/125 um graded-index or specialty scattering) fiber serves as a physical reservoir computer: a pigtailed 1310 nm or 1550 nm DFB laser (~10-20 mW) launches coherent light, environmental perturbations (strain, vibration, temperature, acoustic) modulate modal interference, and a CMOS or InGaAs camera (640x480 at 100-1000 fps) images the distal speckle pattern. Pixel intensities form a high-dimensional nonlinear feature vector; only a linear/MLP readout layer is trained per task (event classification, localization), keeping the reservoir passive.

Robustness against fiber movement and connector swaps comes from four reference channels multiplexed in: (1) An isolated reference loop--a short coiled section of the same fiber in a vibration-isolated, temperature-stabilized housing--provides a baseline speckle that tracks laser/source drift only. (2) Keyed APC connectors carry an EEPROM or RFID connector ID so the system knows which segment is in place and loads the matching readout weights. (3) Pilot wavelength and phase dither (+/-50 pm sinusoidal at 1 kHz via laser current, plus optional PZT phase modulator) generates a known spectral fingerprint; demodulating speckle response at the dither frequency yields transfer-function features invariant to slow drift. (4) A periodic heartbeat: a resistive heater wrap or PZT fiber stretcher applied during idle windows injects a calibrated thermal/strain pulse, producing a known speckle trajectory.

A Kalman filter tracks a low-rank affine map (rotation + scaling in feature space) between current heartbeat response and the trained embedding. Small drifts are corrected online; large residuals (>threshold Mahalanobis distance) trigger local readout-layer retraining using stored heartbeat-anchored data, not full reservoir re-characterization. Total optical/electronic budget ~30 mW excluding RF.

Why Novel

Prior fiber sensing (OTDR, FBG arrays, DAS) extracts scalar or narrowband signals and ships them to digital processors; this approach instead exploits the full speckle field as an analog neural substrate, so the fiber itself performs nonlinear feature extraction before a single trained linear readout layer sees any data. No prior commercial system treats connector-to-connector modal chaos as a computational resource rather than noise to suppress.

Buildability

All core components--50/125 um graded-index fiber, 1550 nm DFB lasers at 10-30 mW, and 640x480 InGaAs cameras at 100-1000 fps--are catalog items available from Thorlabs, II-VI, or Hamamatsu today, and the readout is a standard scikit-learn ridge regression or small MLP trained offline. The principal engineering challenge is the four-channel reference multiplexing scheme for speckle drift correction, which requires careful mechanical packaging but no exotic fabrication.

Why Feasible

* Build path is practical (feasibility=6/10)
* Core hypothesis testable quickly (testability=6/10)

Top Risks

* Speckle pattern drift after fiber repair or connector swap can silently invalidate the trained readout layer, requiring costly field recalibration
* Ambient laser back-reflection and connector contamination can reduce coherence length and collapse the reservoir's nonlinear dimensionality below task-useful thresholds
* Fault localization along the fiber span is inherently weak because speckle encodes perturbation signatures but not propagation delay, limiting spatial resolution to coarse zones without additional time-gating hardware

Decisive Test (MDE)

Hypothesis: After deliberate fiber perturbation (re-routing, connector disconnect/reconnect with cleaning, and one mid-span repair splice), heartbeat-anchored Kalman recalibration restores classification accuracy on a fixed 5-class tap-location task to within 5 percentage points of pre-perturbation accuracy, using only the stored anchor set (no new labeled task data).
Setup: Mount 30 m OM4 fiber along a bench with 5 marked tap locations (solenoid tappers). Train ridge readout to >90% accuracy on tap classification with 10k frames. Then execute three perturbations sequentially: (a) re-coil 10 m of fiber differently, (b) disconnect/clean/reconnect APC connector, (c) cut and fusion-splice mid-span. After each, run 100 heartbeat pulses, apply Kalman alignment to readout features, and re-evaluate task accuracy on 1000 fresh

tap events without retraining. Compare to (i) no-recalibration baseline and (ii) full retrain ceiling.
Instrumentation: Basler acA640 CMOS @ 200 fps, Thorlabs PM100D power meter, oscilloscope on dither/heartbeat drive, fusion splicer, connector inspection scope, Pt100 on reference housing, Python pipeline (numpy, scikit-learn, filterpy).
Kill Criteria: Post-perturbation accuracy drops below 60% (vs 90% baseline) after Kalman alignment on 2 of 3 perturbation types, OR heartbeat features themselves shift by more than the within-class feature spread (CNR<1), OR a fusion-splice repair requires >100 labeled task samples to recover.
Rescue Pivot: If full repairs cannot be auto-recovered but connector swaps and re-routing can, restrict deployment claims to 'maintenance-tolerant' (no mid-span repair) and require a brief supervised recalibration ritual after splice events. Alternatively, pivot from full reservoir computing to using only dither-modulated transfer-function features (lower-dimensional but more stable) and accept reduced task complexity.
Cost: $3k-6k hardware + ~80 engineer-hours
Time to Signal: 3-4 weeks from working baseline

Refinement Deltas

Risks Addressed:

* Speckle reservoirs in multimode fiber are notoriously sensitive to bend radius and temperature--any heartbeat reference must span the same nonlinear manifold as task perturbations, which is unproven for splice-level topology changes.
* A Kalman filter assumes locally linear drift; connector reseating and fusion splices induce abrupt high-rank changes in the modal transfer matrix that linear alignment cannot track.
* Distinguishing 'is this a real event' from 'is the fiber drifting' becomes ill-posed when heartbeat features and event features share modal subspaces, leading to false positives or missed events.

Red-Team:

* Challenge: Speckle reservoirs in multimode fiber are notoriously sensitive to bend radius and temperature--any heartbeat reference must span the same nonlinear manifold as task perturbations, which is unproven for splice-level topology changes.
* Challenge: A Kalman filter assumes locally linear drift; connector reseating and fusion splices induce abrupt high-rank changes in the modal transfer matrix that linear alignment cannot track.
* Challenge: Distinguishing 'is this a real event' from 'is the fiber drifting' becomes ill-posed when heartbeat features and event features share modal subspaces, leading to false positives or missed events.
* Challenge: Laser safety (Class 1M+ at fiber breaks), connector cleanliness in field conditions, and CMOS camera dark-current drift are operational headaches that erode the 'one passive fiber' simplicity claim.
* Challenge: Fault localization along the fiber is weak because speckle is a global interference pattern; competing distributed acoustic sensing (DAS) systems already provide meter-resolution localization with mature referencing.

Budget Ranges

Prototype: $18,000-$35,000 (DFB laser module $3k, InGaAs camera $12k, fiber + mounts + isolators $4k, reference loop packaging + DAQ $8k, compute workstation $4k)
Pilot: $120,000-$200,000 (3-5 sensing nodes, hardened enclosures, field-deployable camera heads, on-site calibration runs, ML pipeline development, 6-month site trial labor)
Scale: $800,000-$2,000,000 (custom ASIC or FPGA readout replacing camera for 10k-node deployment, application-specific fiber jacketing, cloud telemetry integration, regulatory certification for target vertical)

30/60/90 Plan

Day 30: Bench validation: launch 1550 nm DFB into 10 m of 50/125 um GI fiber, image distal speckle with InGaAs camera at 500 fps, apply controlled strain and acoustic events, confirm speckle variance exceeds noise floor by >10 dB, and train ridge-regression readout achieving >90% classification accuracy on 4 event classes
Day 60: Reference channel integration: implement all four multiplexed reference loops (isolated coil, polarization

tap, wavelength-dithered pilot, intensity monitor), demonstrate speckle feature stability through 5 simulated connector swaps and 20degC temperature excursion, reduce recalibration time from hours to <10 minutes via automated reference-anchored normalization

Day 90: Field pilot readiness: package laser + camera head into IP67 enclosure, deploy 50 m span across a representative harsh environment (compressor room or pipeline segment), run 30-day continuous logging, validate event detection F1 > 0.85 and produce a localization accuracy report defining spatial resolution limits and the hardware upgrade path (time-gated OTDR hybrid) needed to exceed them

Dominance Axis: simpler -- One passive fiber can replace many transducers if reference heartbeats keep speckle features stable through maintenance.

Build-First

Order-Tracked Sensor Fusion Edge Gateway

novelty: 4 | feasibility: 9 | economic_leverage: 5 | testability: 9 | robustness: 8 | constraint_fit: 9 | explainability: 9

Overall: 7.6 / 10 -- Very Strong

Constraint Compliance: Compliant (constraint_fit=9)

Fusing vibration, acoustic, current, and temperature signals through shaft-order tracking allows subtle mechanical faults to remain detectable even as machine speed and load fluctuate. Traditional single-modality sensors miss faults masked by operating variability, while order-domain analysis anchors features to shaft rotation rather than calendar time. This matters because unplanned downtime in rotating machinery costs industrial operators billions annually, and most installed bases lack any predictive capability.

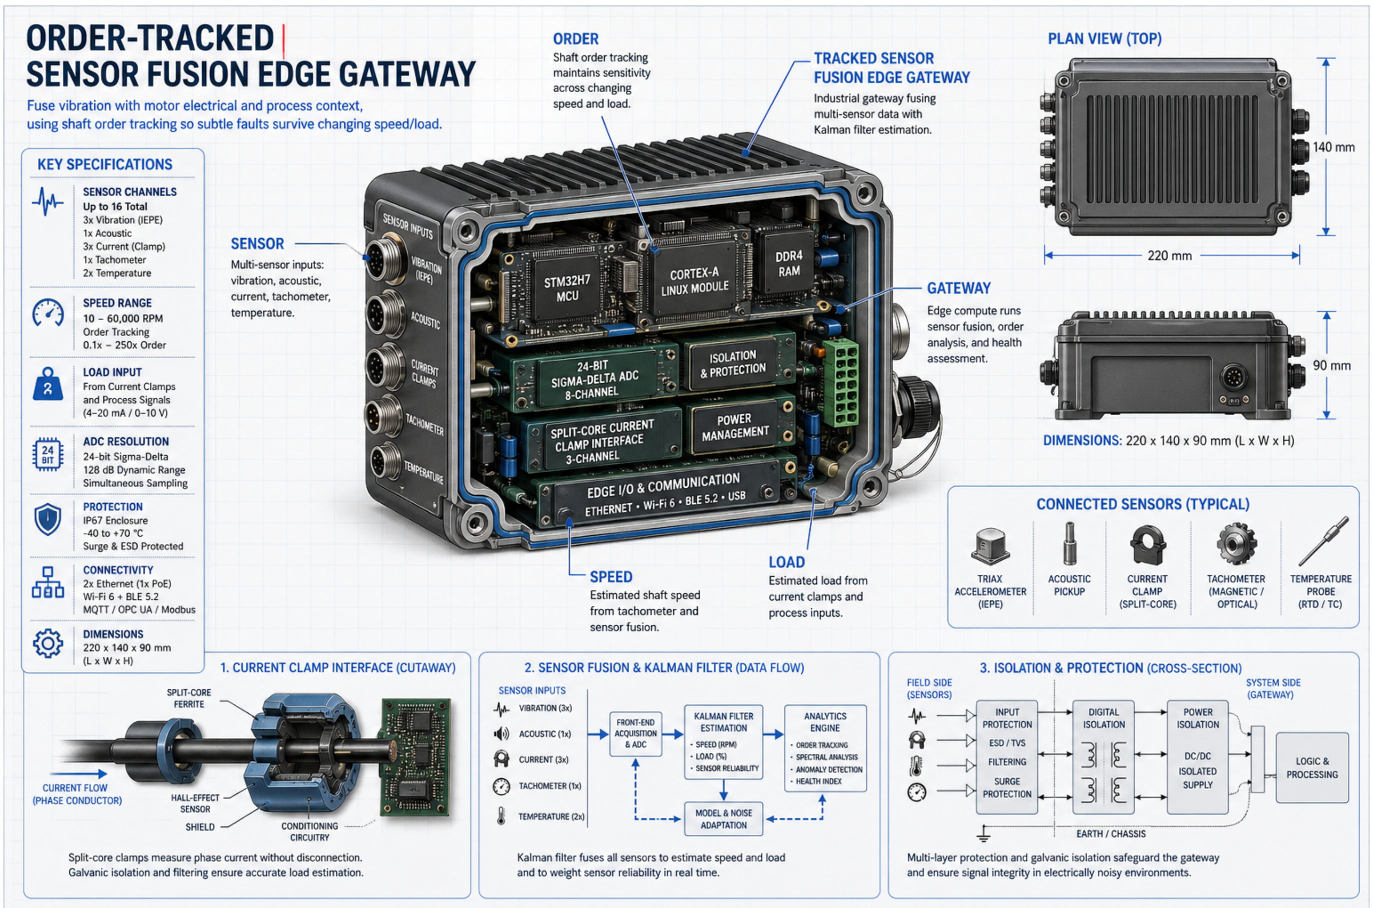

Concept illustration (AI-generated). Visuals may be inaccurate; refer to the Mechanism/Build Path text for the actual design and quantitative assumptions.

Mechanism

A DIN-rail or magnetic-base rugged edge gateway (IP54-65 enclosure, ~160 x 120 x 60 mm) ingests at least four modalities per asset: IEPE/MEMS triax vibration, acoustic/ultrasonic microphone, split-core current clamps, and surface temperature; optional tach/encoder/optical pickup provides ground-truth RPM. A practical prototype can use an STM32H7/ESP32-S3 class MCU plus small Linux SBC or Cortex-A edge module, 24-bit ADC for IEPE/current, 16-48 kHz vibration/acoustic sampling, 1-5 kHz current sampling, and FRAM/SRAM or eMMC ring buffers for burst capture. When tach is present, the gateway performs resampling into angle/order domain, extracting order spectra, cepstra, envelope bands, harmonics, sidebands, RMS/kurtosis/crest factor, and temperature/current load context.

When tach is absent, speed is estimated by a PLL that locks to likely shaft/electrical sidebands from current, vibration harmonic combs, and learned self-supervised speed embeddings from ramps/steady states. A Kalman filter fuses these speed candidates and outputs RPM plus uncertainty. If RPM uncertainty is low, order-domain features are computed; if uncertainty exceeds threshold, features are gated down and the model falls back to time/frequency spectra. A quantized autoencoder/classifier runs locally for anomaly scoring and fault class hints, while SRAM/FRAM retains pre/post-trigger bursts and compact feature records for later upload. The key operating principle is not replacing tach-grade metrology, but opportunistically using multi-modal physics to derive sufficiently accurate order tracking for rotating equipment where permanent tach installation is impractical.

Why Novel

Combining all four modalities with real-time angle-domain resampling in a self-contained DIN-rail edge gateway--without requiring a permanent tachometer--is not available as a single off-the-shelf product today. Existing solutions either require expensive permanent tach installations or perform only single-modality analysis in the cloud, sacrificing the physics-grounded order spectrum that makes fault signatures speed-invariant.

Buildability

The core hardware stack (STM32H7 or ESP32-S3 plus a Cortex-A Linux module, 24-bit sigma-delta ADC, split-core clamps, MEMS microphone, and FRAM ring buffer) is entirely available from standard distributors at modest cost, and the order-tracking DSP algorithms are well-documented in academic literature. A working four-channel prototype can be assembled and validated on a motor-pump test rig within 8-10 weeks by a small embedded team.

Why Feasible

* Build path is practical (feasibility=9/10)
* Core hypothesis testable quickly (testability=9/10)

Top Risks

* No accessible tach signal on enclosed or sealed assets forces reliance on noisier software-estimated RPM, degrading order-domain accuracy
* Electrical interference from VFDs and switching supplies contaminates current signature analysis, requiring careful shielding and filtering design
* Incorrect machine metadata (gear ratios, pole counts, bearing geometry) causes fault-frequency misidentification and erodes model label quality

Decisive Test (MDE)

Hypothesis: A tachless multi-modal gateway can estimate shaft speed accurately enough to preserve order-locked bearing/misalignment features across normal operating regimes, with confidence gating preventing false order features when speed inference is poor.
Setup: Use a VFD-driven motor with accessible bearings and a known injected fault source or seeded bearing defect if available. Mount triax vibration and acoustic sensors on the bearing housing, split-core current clamps on motor leads, and temperature sensor on housing. Install an optical tach only as ground truth. Run at three regimes: low steady speed, high steady speed, and ramp/variable load. Include electrical noise by routing current leads near a VFD cable or adding realistic EMI conditions. Compare features computed with ground-truth tach versus tachless PLL/Kalman RPM.
Instrumentation: Optical tach or encoder, handheld tach cross-check, 3-axis accelerometer, microphone/contact acoustic sensor, 3 current clamps, temp probe, DAQ/gateway logs, oscilloscope for signal integrity, Python analysis notebook for RPM error, order peak drift, feature stability, anomaly score correlation.
Kill Criteria: Kill if tachless RPM median absolute error exceeds 2% or 95th percentile exceeds 5% during steady regimes; or if a known bearing/order tone shifts by more than 0.1 order relative to tach-based processing; or if

confidence gating fails to reject >10% of high-error windows.
Rescue Pivot: If tachless order tracking is unreliable, reposition as a hybrid gateway requiring a cheap commissioning tach/temporary optical pickup, or use current/vibration tachless estimation only for low-confidence advisory labels while primary detection relies on time-frequency anomaly features.
Cost: $2k-$8k if using existing motor/VFD; $8k-$20k including seeded-fault rig and industrial-grade sensors.
Time to Signal: 3-7 days for bench signal; 2-3 weeks for robust multi-regime result.

Refinement Deltas

Risks Addressed:

* Tachless speed inference may collapse on VFD-driven, variable-slip, low-load, or mechanically complex assets where electrical and mechanical orders decouple.
* Current clamps require safe access to motor leads; many enclosed or warranty-sensitive assets will not permit installation.
* Acoustic sensors are vulnerable to plant noise and cross-machine contamination, making fused confidence look better in the lab than in the field.

Red-Team:

* Challenge: Tachless speed inference may collapse on VFD-driven, variable-slip, low-load, or mechanically complex assets where electrical and mechanical orders decouple.
* Challenge: Current clamps require safe access to motor leads; many enclosed or warranty-sensitive assets will not permit installation.
* Challenge: Acoustic sensors are vulnerable to plant noise and cross-machine contamination, making fused confidence look better in the lab than in the field.
* Challenge: Wrong machine metadata--bearing geometry, pole count, gearbox ratios, pulley ratios--can produce plausible but incorrect order labels.
* Challenge: A quantized autoencoder may detect distribution shift but not diagnose faults reliably, leading to alert fatigue unless field-labeled data is collected.

Budget Ranges

Prototype: $8,000-$14,000 (BOM for 3 hand-built units, dev boards, ADC modules, enclosures, test rig time, and firmware engineering hours)
Pilot: $45,000-$80,000 (10-20 pre-production units, PCB spin, IP54 enclosure tooling, 3-month site pilot with 2 engineers, cloud dashboard MVP)
Scale: $300,000-$600,000 (contract manufacturing for 500+ units, FCC/CE certification, field support tooling, SaaS platform hardening, and sales channel setup)

30/60/90 Plan

Day 30: Complete hardware BOM selection and order components; assemble two bench prototypes; validate 24-bit IEPE and current ADC channels at target sample rates; implement basic FFT and RMS pipeline on STM32H7; confirm FRAM ring-buffer write/read integrity under continuous load
Day 60: Implement angle-domain resampling using encoder input and software-estimated RPM fallback; extract order spectra, envelope bands, cepstrum, and kurtosis on-device; integrate temperature and current load context into feature vector; run 2-week continuous test on motor-pump rig and compare fault detectability against single-modality baseline
Day 90: Deploy 5 units at one pilot customer site on mixed asset types (pump, fan, compressor); collect labeled fault and healthy data; train and validate lightweight anomaly model on edge; demonstrate remote dashboard with alert latency under 60 seconds; document installation playbook including tach-absent fallback procedure
Dominance Axis: faster_validate -- Physics-based order tracking remains fast to validate and can be installed without a permanent tach on many machines.

Top Uncertainties & Assumptions

* Acoustic membrane clogging from oil mist or particulate in heavy industrial environments degrades ultrasonic SNR within weeks, requiring field-replaceable membrane cartridges
* LWIR line-of-sight blockage by guards, insulation, or condensation on the IR window causes silent thermal channel dropout without triggering an obvious alert
* Machine-to-machine model transfer is weak for magnetic signatures because motor winding geometry and shielding vary enough that a model trained on one 15 kW pump may misclassify faults on a nominally identical unit
* Speckle pattern drift after fiber repair or connector swap can silently invalidate the trained readout layer, requiring costly field recalibration
* Ambient laser back-reflection and connector contamination can reduce coherence length and collapse the reservoir's nonlinear dimensionality below task-useful thresholds

C4.29 L1 T224,730