



OPERATIONS INTELLIGENCE

Catch Problems Before They Cost You

Operators are drowning in alerts



Industrial monitoring systems generate thousands of alerts every day. But when 70–80% of them are false, operations teams lose the ability to distinguish routine sensor noise from early signs of real equipment failure.

Critical degradation — a compressor vibration trend, a pressure drift in a separator — goes unnoticed among hundreds of irrelevant notifications. By the time the problem becomes obvious, it's already an emergency

ALERT OVERLOAD

70 - 80% of alerts are false positives

Field teams spend their shift chasing sensor noise. Real equipment degradation is buried under thousands of irrelevant notifications

ALARM FATIGUE

Operators stop trusting the system

When every alert looks the same, crews dismiss warnings as noise. A compressor drifting towards a trip? Another false alarm. Until it isn't

REACTIVE MAINTENANCE

5 - 10x higher cost than planned repairs

Equipment fails without warning. Emergency mobilization offshore turns a routine repair into a six-figure event

PRODUCTION SHUTDOWN

Weeks to months of lost output

A missed compressor drift or well integrity issue can shut down an entire FPSO. Deferred production alone costs millions – before you count repairs

Every missed anomaly has **A PRICE TAG**

The financial exposure from undetected equipment failures in upstream operations is well documented — but rarely visible until it's too late. A single missed vibration drift on an offshore compressor can escalate from a routine inspection to a full production shutdown within days. The longer degradation goes unnoticed, the more expensive the recovery

SCENARIO 1

Unplanned compressor trip on an offshore platform

\$1 – 3M per day

in deferred production, emergency logistics, and accelerated part procurement

SCENARIO 2

Missed well integrity anomaly requiring emergency intervention

\$10 – 50M+

in well workover, environmental remediation, regulatory investigation, and potential production curtailment

SCENARIO 3

Shutdown due to cascading equipment failures across multiple systems

Months of downtime

can cause full production halts, fleet mobilization, and regulatory review – the costliest outcome in offshore operations

From alert noise to actionable intelligence

ReviveAI connects directly to plant historians, SCADA, and DCS systems already deployed on site – without replacing or reconfiguring existing infrastructure. The platform ingests live sensor streams, **applies multiple independent detection algorithms simultaneously**, and cross-validates results to separate out genuine equipment degradation from routine sensor noise. What reaches the operations team is not another flood of raw alerts, but **a prioritized queue of confirmed events**, each accompanied by equipment context, maintenance history, a confidence score, and a ready-to-assign work order

NOW

Thousands of raw alerts per shift with no prioritization – operators decide what matters based on gut feel



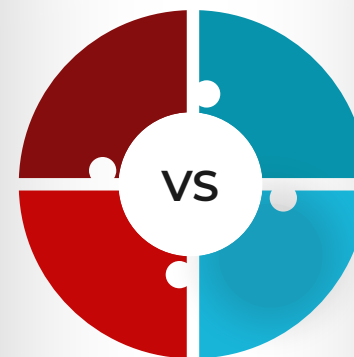
Single-threshold alarms that trigger on any deviation, producing up to 80% false positives



Manual work order creation – by the time paperwork is done, the shift is over and context is lost



No connection between the alert and the equipment history – each event is evaluated in isolation



REVIVE AI



Events ranked by severity and confidence – critical items surface first, low-confidence noise is filtered out automatically



Multi-algorithm consensus – an anomaly is only escalated when 2 or more independent detection layers agree



Pre-drafted work orders with full evidence – sensor data, expected vs. actual values, maintenance record, and a shift handover note



Each event linked to equipment context – asset type, location, risk category, last service date, and open work orders

Designed for Industrial Operations



ReviveAI adapts to the operational context of each team — whether the focus is uptime, commissioning quality, or long-term asset performance



Plant Managers & Reliability Engineers

Owner-operated facilities, JV partners

Responsible for uptime and asset health across dozens of rotating and static equipment units. Current alarm systems flood the control room with undifferentiated notifications, making it nearly impossible to separate early degradation signals from background noise — especially across multiple shifts

What ReviveAI delivers

- Early warnings on process deviations — days or weeks before a threshold alarm would trigger
- Maintenance prioritization by risk — critical events surface first, with evidence and recommended actions
- Shift handover context — every event carries a summary that the next crew can act on immediately



EPC Contractors & Commissioning Teams

Engineering firms, system integrators

During project handover and commissioning, teams need to verify that instruments are reading correctly, process flows match design specifications, and all numbers fall within acceptable ranges. Doing this manually across hundreds of sensors is slow and error-prone

What ReviveAI delivers

- Automated validation of instrument readings against expected operating envelopes during commissioning
- Deviation flagging against design constraints — catch scope discrepancies before they become an issue
- Auditable evidence trail — every reading, comparison, and flag is documented for handover sign-off



Owner-Operators & Asset Management


NOCs, IOCs, independent producers

Managing long-lifecycle assets across multiple fields requires continuous visibility into equipment condition — not periodic inspections. Existing CMMS and historian systems hold valuable data, but lack the intelligence layer to connect sensor trends with maintenance decisions at scale

What ReviveAI delivers

- Continuous anomaly detection across all connected assets — not sampling, not periodic, always running
- Auto-generated CMMS-ready work orders — with evidence, priority, and suggested owner pre-filled
- Full operational traceability — from sensor event to work order to resolution, all linked and auditable

Our Technologies Powers



Architecture Overview



Operations Dashboard – full situational awareness



The dashboard displays the following information:

- Active wells:** 1 AI across 1 fields & terminals
- Anomalies today:** 200 (+149.1% vs 7-day avg comparison)
- Critical alerts:** 0 requiring intervention
- Production at risk:** 21060k of 135k bopd total
- MTBF (rotating equipment):** 142 days mean time between failures
- HSE status:** Green (0 open MUPRC items)
- Open work orders:** 0 across SAP PM
- All Assets Configured:** All 1 assets have complete configuration data.
- Asset Map:** Geospatial view of Nigeria with a highlighted asset location.
- Live Anomalies:** 100 active anomalies. Filtered by severity (All, Critical, High, Medium, Low). Includes escalation routes (MS Teams - Lagos Ops, SMS - On-call, Email - HSE).
- Anomaly Details:**
 - Temperature_Heat_Exc...** (HIGH): Temperature_Heat_Excchanger_1 ML algorithms detected... Value: 70%
 - Pressure_Wellhead_2** (MEDIUM): Pressure_Wellhead_2: ML algorithms detected... Value: 14%

KPI PANEL

Production exposure at a glance



Active wells, anomaly count with trends, critical alerts, and barrels per day at risk – the four numbers that define operational risk posture right now

HEALTH METRICS

MTBF, HSE status, open work orders



Rotating equipment reliability, regulatory compliance status, and maintenance backlog – visible without opening SAP or a separate reporting tool

ASSET MAP

Geospatial view of all monitored assets



Every connected field, terminal, and FPSO plotted on the map with real-time health status. Drill into any asset for sensor-level detail

LIVE ANOMALIES

Real-time alert feed with severity and scores



Each anomaly shows the sensor tag, current value, confidence score, and severity level. Operators acknowledge or escalate directly from the feed

Operations Dashboard – full situational awareness



FULL DEVICE GRID

All sensors visible at once



No clicking through menus to find a tag. Every sensor in the connected stream is displayed as a card with its sparkline, operating range, and current reading – sorted by anomaly count so the noisiest tags surface first

ANOMALY DENSITY

Spot problem clusters instantly



The badge on each card shows how many anomalies that device has produced. When multiple related devices light up simultaneously, it points to a systemic issue rather than isolated sensor noise

CROSS-CORRELATION

From individual tags to system behavior



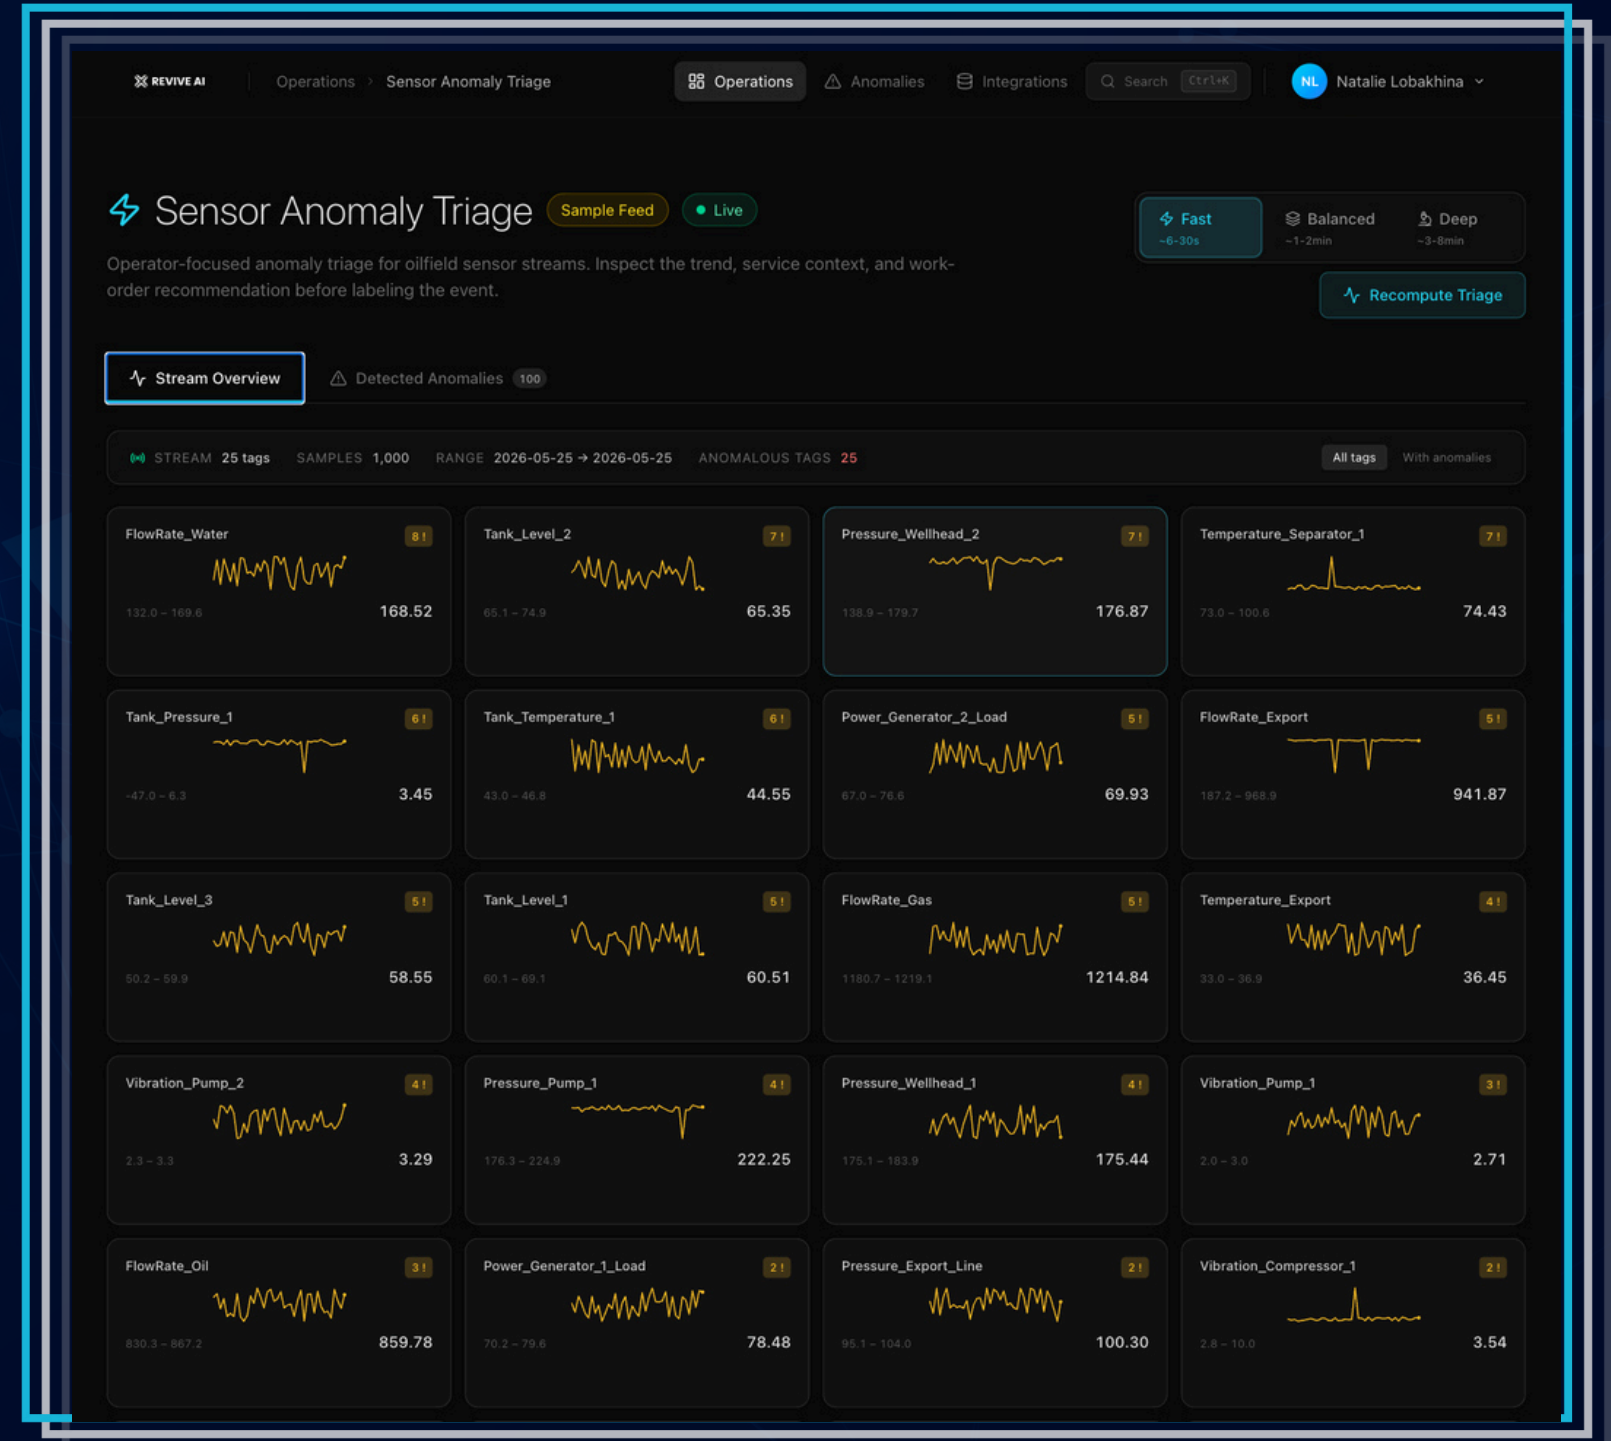
When multiple sensors all show elevated anomaly counts for the same time window, the pattern tells a story that no single sensor could. The grid makes these correlations visible without running a separate analysis

TRIAGE MODES

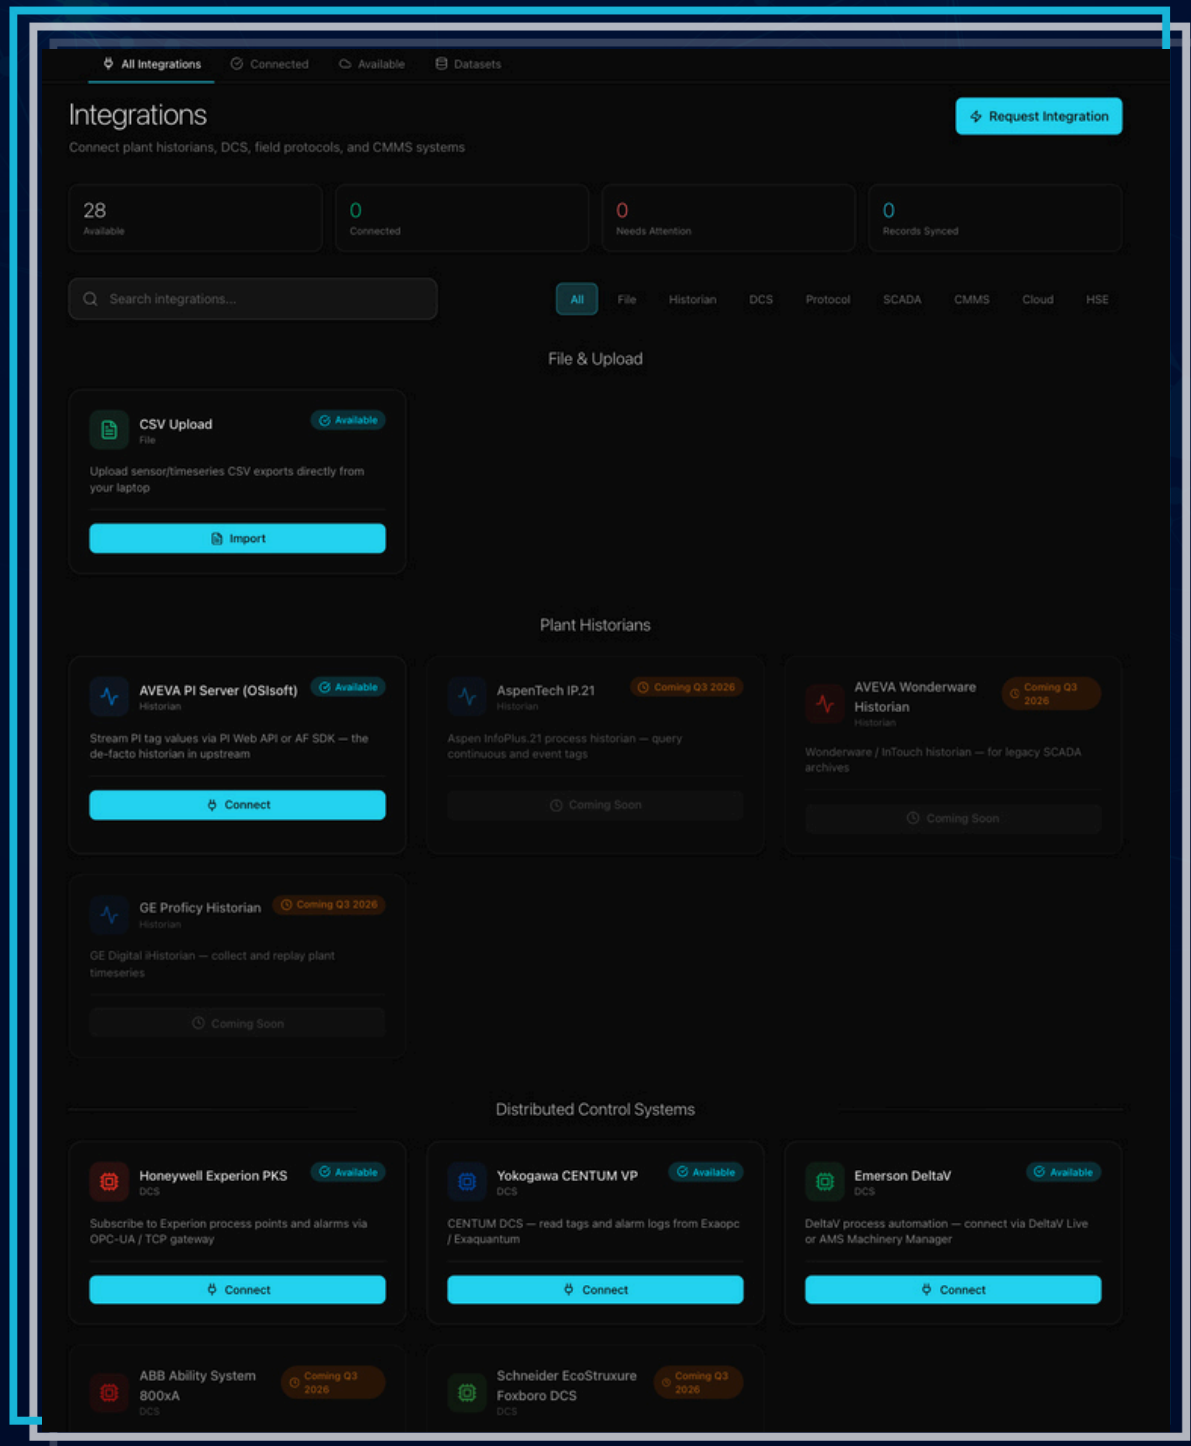
Fast, Balanced, or Deep analysis



Run a quick 6–30 second scan for shift handover, or a deep 3–8 minute analysis for investigation. The same stream, different levels of scrutiny – chosen by the operator



Integration Ecosystem



ReviveAI is not a rip-and-replace platform. It plugs into the data infrastructure that upstream operators have already invested in – plant historians, distributed control systems, SCADA networks, and CMMS platforms. No new hardware, no middleware, no reconfiguration of existing tag databases

28 integrations across 8 categories

File upload, plant historians, DCS, field protocols, SCADA, CMMS, cloud services, and HSE systems – searchable and filterable from a single console

The systems that hold the data

Direct connectors to the historian platforms deployed across 90%+ of upstream operations worldwide

Real-time process control data

Subscribe to process points and alarms directly from the control system – no export step, no data loss

CSV upload for day-one evaluation

No integration project required to get started. Export sensor data from any historian as CSV, upload it, and run anomaly triage immediately. The live connector can be configured in parallel

Anomaly detection workflow **Demo**



Sensor Anomaly Operations

ReviveAI

Subsea Pump P-3 Suction Pressure (bar)

Risk: Subsea pressure delivery anomaly

2 LAYERS POINT Wavelet Transform Seasonal Hybrid ESD

Dec 24, 2025 +2.7%

Cargo Tank T-05 4x TIMES

Storage systems - TANK T-05 LEVEL (m)

Risk: Tank level/pressure control anomaly

2 LAYERS POINT IQR Method Seasonal Hybrid ESD

Dec 22, 2025 -3.8%

Crude Meter FT-901 4x TIMES

Export metering - CRUDE METER FT-901 DENSITY (kg/m3)

Risk: Meter drift / allocation quality risk

3 LAYERS POINT Robust Z-Score IQR Method +1 more

Dec 13, 2025 +0.3%

SUBSEA RISER R-1 TOP 4x TIMES

SCADA sensor stream - SUBSEA RISER R-1 TOP PRESSURE (bar)

Risk: Sensor trend outside learned operating envelope

2 LAYERS POINT IQR Method Seasonal Hybrid ESD

Dec 8, 2025 -2.8%

Provide Feedback (Learning System)

Help improve detection accuracy by labeling this anomaly:

Reason Code

Equipment issue Phase imbalance Sensor issue Seasonal load Scheduled maintenance Normal operation

Data artifact

Operator Notes (Optional - Learn Context)

Example: "Scheduled maintenance was performed yesterday, this is expected behavior" or "Confirmed with field crew - found broken insulator on phase B"

Tip: Detailed notes improve AI accuracy by 40-50% after just 300 labeled examples.

Apply learning to similar anomalies

Use this feedback as a signal for pending anomalies on the same metric.

Impact Preview: Real Issue
This will confirm 1 anomaly and keep it high priority.

Impact Preview: Normal
This will mark 1 anomaly as normal and lower its priority.

REAL ISSUE **NORMAL**

Watch on YouTube

Feedback History

No feedback history for this anomaly yet.

Why Operations Teams Choose Us



Intelligent Sensor Data Ingest

Turn SCADA streams, historians, CSV exports, and field tags into a unified, structured monitoring layer

- ✓ Sensor coverage across process, mechanical, electrical, instrumentation, and packaged systems
- ✓ Tag metadata, units, asset hierarchy, and quality flags normalized automatically
- ✓ Faster alignment between operations, reliability, and engineering teams

Multi-Algorithm Consensus Detection

Align every detection with 13 proven algorithms, asset baselines, and your team's prior judgment

- ✓ Consensus across point, drift, regime-change, seasonal, wavelet, and oscillation detectors
- ✓ Expose subtle anomalies, sensor drift, and silent failure modes
- ✓ Standardize alert decisions across control room, reliability, and field teams

Downtime & Risk Visibility

Surface the issues that erode uptime, safety, and production confidence – before they escalate

- ✓ Early warning on drift, regime shifts, and degrading equipment before failure
- ✓ Cross-tag correlation, cascade risk, and unplanned-shutdown exposure
- ✓ Severity tiers, confidence scores, and operator-tuned false-positive suppression

Automated Insight & Reporting

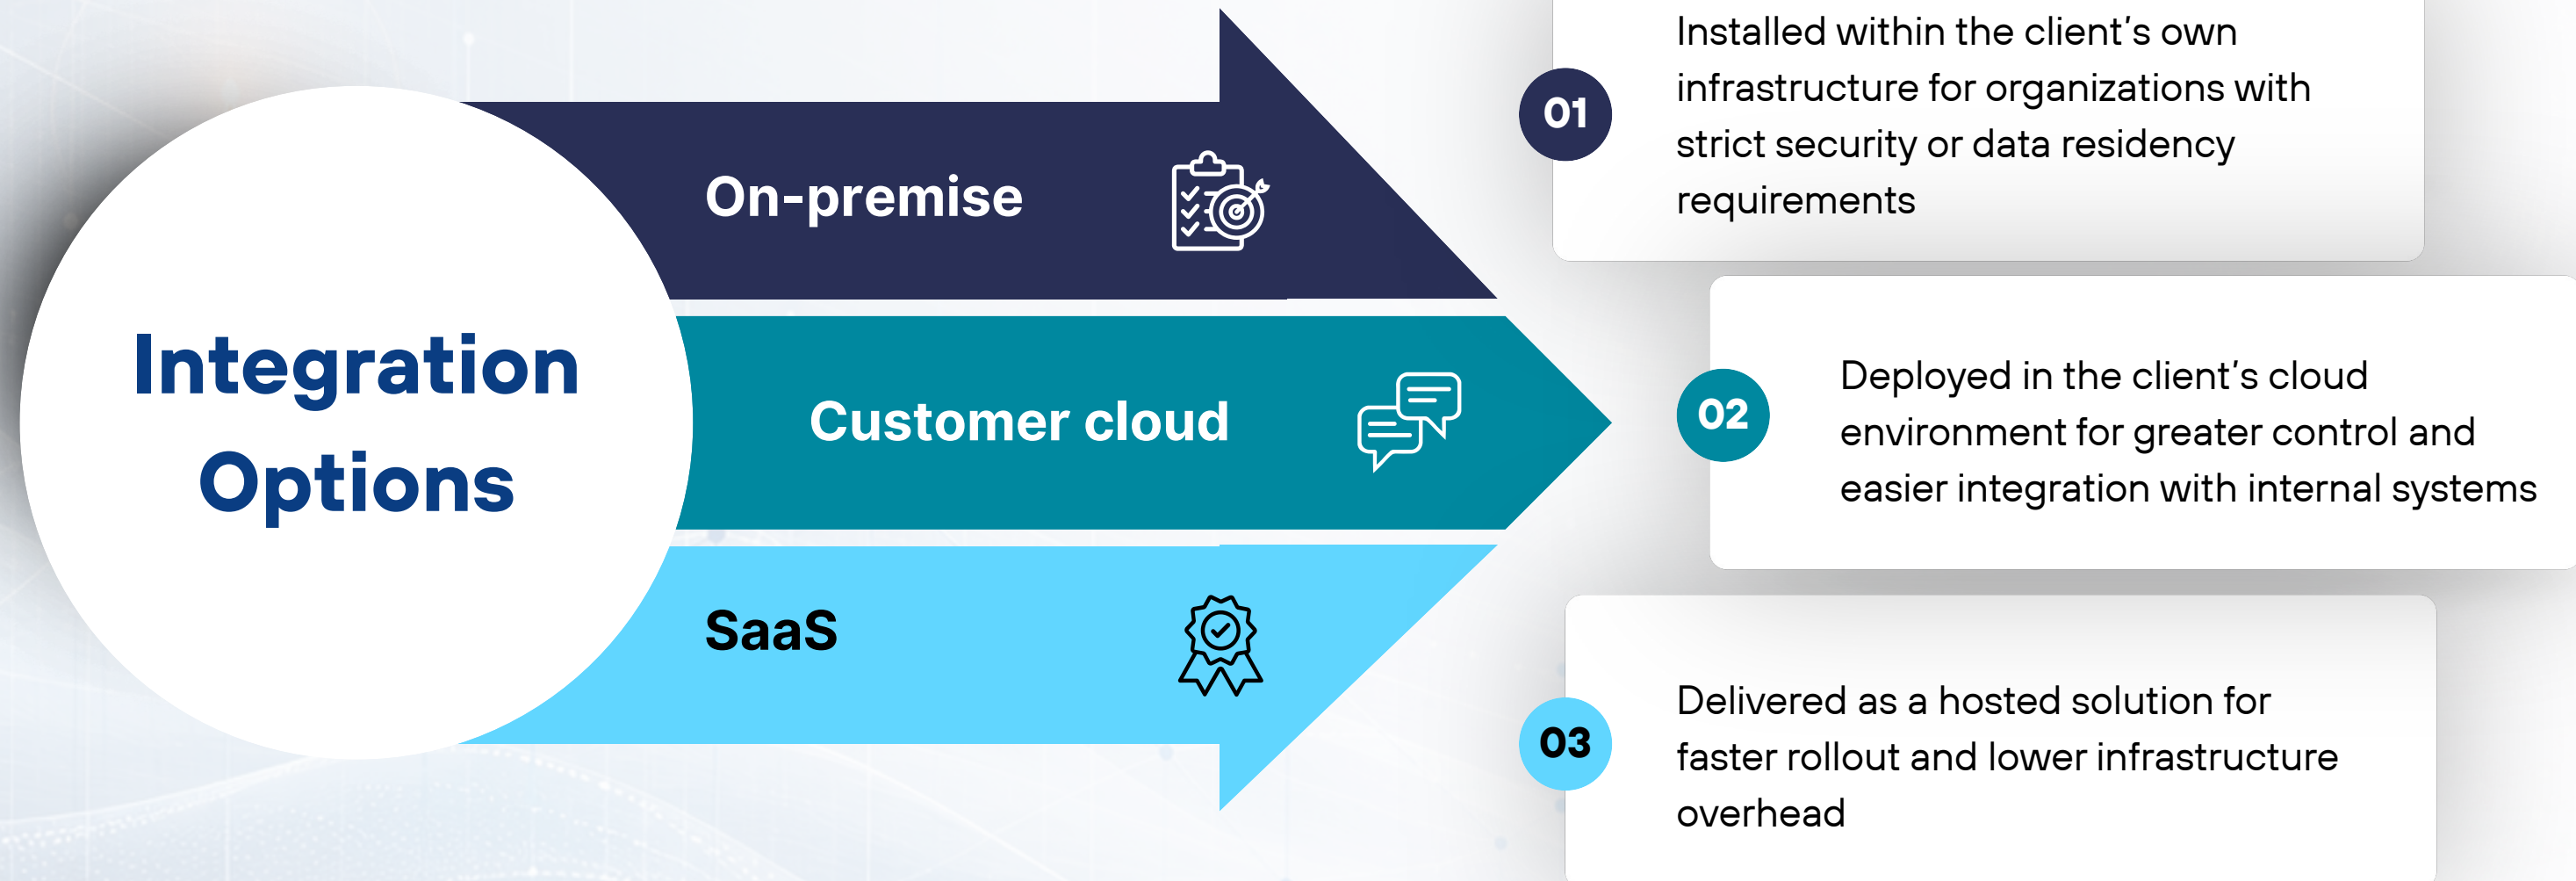
Generate faster, more consistent operational reports for control room, leadership, and audit

- ✓ Anomaly summaries and incident timelines per asset
- ✓ Operator feedback logs and detection-accuracy dashboards
- ✓ Shift handover briefs, KPI rollups, and executive-ready outputs

Deployment & Integration Options



Revive AI Operations Intelligence can be deployed in the model that best fits the client's IT, security, and compliance requirements:



About us



Revive AI started with a question: how do you build AI that never hallucinates?

LLM-Assisted Translation and Bounded Model Checking of Python Code

Shivkumar Shivaji¹, Klaus Havelund², Natalia Lobakhina¹, Lucas C. Cordeiro^{3,4}, and Alessandro Pinto²

¹ Generative AI LLC, USA

shiv@revive.cx, nlobakhina@revive.cx

² NASA Jet Propulsion Laboratory, California Institute of Technology, USA
klaus.havelund@jpl.nasa.gov, alessandro.pinto@jpl.nasa.gov

³ University of Manchester, UK

lucas.cordeiro@manchester.ac.uk

⁴ Federal University of Amazonas, Brazil

Abstract. Formal verification of Python programs remains challenging due to the language's dynamic nature and rich semantic constructs. We present an approach that combines Large Language Models (LLMs) with Bounded Model Checking (BMC) to verify Python code. Our system uses an LLM to translate Python programs into C code suitable for formal verification using ESBMC, enabling the detection of critical bugs, including arithmetic overflows, array bounds violations, and concurrency errors. We evaluate our approach on a benchmark of 23 carefully designed Python programs with planted bugs representing common verification challenges. The LLM orchestrator successfully detected all planted bugs by iteratively coordinating static, dynamic, and formal verification tools. Our results demonstrate that LLM-assisted translation can make mature C verification tools accessible for Python code analysis, though the current evaluation is scoped to programs of 15–50 lines with bounded loops and statically-sized data structures, a practical limit driven by ESBMC verification times rather than a fundamental BMC constraint.

Keywords: Formal Verification · Bounded Model Checking · Python · Large Language Models · Code Translation

We found the answer in formal verification – mathematical methods used to prove software correctness in aerospace and safety-critical systems.

Our research, published with NASA scientists, showed how AI can assist in verifying complex code. Then we asked: where else do people need verified accuracy? Where does one missed signal cost millions?

Now we apply these same verification principles to real industrial operations: monitoring SCADA streams and historian data across wells, separators, compressors, and substations — where a false alarm wastes an operator's shift, and a missed anomaly becomes an unplanned shutdown costing millions per day.

We're building AI for the industries where mistakes are expensive – and quality is non-negotiable.

Founding Team

Our team is combining deep expertise in AI, infrastructure, and product development with a proven track record of delivering impactful solutions



Shivkumar Shivaji

- **Ph.D. in Computer Science** from the University of California, Santa Cruz
- **25+ Years of Expertise:** Extensive experience in software development, machine learning, and Agentic AI.

Proven Innovator:

- Built a machine learning platform for predicting and analyzing storage performance at Western Digital.
- Developed an Agentic AI solution for Fortune 400 companies to enhance code quality and security through automated unit test generation.
- Holds a patent for storage anomaly detection (Patent No. 10372524), showcasing expertise in efficient monitoring of large-scale storage systems.

Co-founder and CTO in:

- Crowdscience (acquired by Yume)
- PatternEQ (acquired by Western Digital)
- AutonomIQ (acquired by Sauce Labs)



Natalia Lobakhina

- **PhD in Economics, Master of Psychology**
- **15+ Years of Expertise** in entrepreneurship, business consulting, and technology startup acceleration

Proven Expertise:

- Launched and accelerated over 70 IT projects, securing funding for 20+ startups from commercial and government investors.
- Developed and implemented methodologies for assessing innovative potential and startup sustainability
- Lead consultant for business valuation, sales development, and growth strategy for fast-scaling companies and startups.

Specializations:

- Startup acceleration, technology commercialization, business investment and scaling, innovation management, strategic development, team and talent development.

Contact Us:

Shivkumar Shivaji
Natalia Lobakhina

SCHEDULE TECHNICAL DEMO

(510) 304-1184



Palo Alto, California



Shiv@revive.cx



www.revive.cx



 **REVIVE AI**

