

EdgeAlyzer: True Edge AI



EdgeAlyzer offers true predictive maintenance AI



Utilize EdgeAlyzer without cloud subscriptions



Eliminate IT headaches with our edge solution



EdgeAlyzer: Smart Box Explained

Core Capabilities



A complete, self-contained predictive maintenance platform



Provides true AI functionality directly at the edge



Turnkey IoT solution including sensors and analytics

Implementation & Benefits



Built by experienced industrial automation engineers



Requires zero IT infrastructure and has no subscriptions

The Business Problem We Solve

Unplanned Downtime is the Enemy

The Reality:



• A single hour of unplanned downtime can cost **\$1,000 to \$250,000+** depending on the operation



- Traditional maintenance is either **reactive** (fix it when it breaks) or **preventive** (replace parts on a schedule whether they need it or not)
- Reactive = expensive emergency repairs and lost production
- Preventive = wasted money replacing parts that still have life left

What Predictive Maintenance Promises:

Know **when** equipment will fail, **before** it happens — so you can schedule repairs during planned shutdowns.

Why Most Attempts Fail:

Approach	The Promise	The Reality
Cloud Analytics	Upload your data, we'll find problems	IT barriers, latency, subscription costs, data leaves facility
Vibration Consultants	We'll analyze your equipment monthly	Expensive, infrequent, findings arrive too late
"Smart" Sensors	AI-powered monitoring	Usually just simple limits with marketing buzzwords

How EdgeAlyzer Solves It:

✓ A **complete, self-contained system** that:

- ✓ Monitors continuously (not monthly snapshots)
- ✓ Analyzes instantly (not uploaded for later processing)

✓ Learns **YOUR** equipment (not generic patterns)

✓ Requires no IT infrastructure (plug in and go)

Result: Real predictive maintenance that actually works — on your timeline, under your control.

Why Real-Time Matters – The Data Fidelity Problem

The Difference Between Historical Analysis and Live Inference

Most "AI" monitoring systems work like this:

- 1. Sensors collect data and store it
- 2. Data is batched and uploaded to the cloud (hours or days later)
- 3. Cloud servers run analysis on historical data
- 4. Results are sent back (eventually)

The Problem:

By the time you get an answer, the moment has passed.

Physical Fidelity – Why Timing Matters:

What's Happening	What Data Collection Captures	What You Miss
Motor starts under load	Samples taken every few minutes	The startup vibration spike that reveals bearing wear
Spindle accelerates through resonance	Averaged values uploaded hourly	The momentary frequency that indicates imbalance
Pump cavitates briefly	Batch upload next morning	The pressure/vibration correlation that predicts seal failure

EdgeAlyzer's Approach: **Collect** → **Analyze** → **Decide** – All in milliseconds, on the same device.

- ✔ **No batching** – every reading is analyzed as it arrives
- ✔ **No upload delays** – the intelligence lives on the machine
- ✔ **Immediate response** – alerts happen when problems happen, not hours later
- ✔ **No lost context** – speed, load, and vibration are correlated in real time

The machine's behavior right now tells you more than terabytes of historical data ever could.

From Data Mountains to Actionable Signals

The Dirty Secret of Industrial IoT

We've seen it firsthand:

- ✗ • Terabytes of sensor data stored in databases
- ✗ • Beautiful dashboards showing historical trends
- ✗ • Zero actionable alerts ever generated
- ✗ • Equipment still fails unexpectedly

The Data Hoarding Trap:

What Companies Do	What They Get
Collect everything "just in case"	Storage costs, no insights
Build dashboards to visualize data	Pretty pictures, no action
Hire data scientists to find patterns	Months of analysis, outdated by the time it's done
Wait for someone to notice a problem	Reactive maintenance disguised as "monitoring"

The Missing Link: Decision at the Point of Collection

Data has no value until it becomes a decision.

EdgeAlyzer doesn't just collect data – it **converts every reading into a yes/no answer**:

Input	EdgeAlyzer Decision	Output
Vibration + Speed + Current	Normal for this operating state?	✅ Continue monitoring
Vibration + Speed + Current	Unusual pattern detected?	⚠️ Alert maintenance now

The Result:

- ✅ **No data scientists required** – the system makes the call
- ✅ **No 'check the dashboard' culture** – alerts come to you
- ✅ **No analysis paralysis** – clear signals, not overwhelming data
- ✅ **No terabytes of useless logs** – store what matters, act on what's urgent

EdgeAlyzer turns sensor readings into maintenance work orders – automatically, in real time.

How is it **Installed?**



No complicated data management. No custom code. No IT involvement.

Problem 1 – IoT Installation is Always an Engineering Project

The Issue:



- Traditional platforms and sensors **sold separately** – custom integration required



- Connecting to PLCs requires Modbus/OPC-UA configuration on both sides



- MTConnect is “standard” but implementations vary – parsing XML is custom for each machine



- Additional sensors = additional PLCs on customer network = permanent IT burden

The EdgeAlyzer Solution:



- Sensors are **isolated on a local network** – doesn't touch customer infrastructure



- Everything **preconfigured**: 40 vibration + 40 analog + 40 digital I/O – setup via simple menus



- **No custom code. No PLC modifications. No XML parsing.**

Problem 2 – Traditional Systems Rely on Customer IT

The Issue:



- Weeks lost to compatibility issues with security setups, VMs, network configurations



- Running systems taken down by customer updates, maintenance, security patches



- IP conflicts crash systems (surprisingly common in industrial environments)



- STIG and security reviews block integration

The EdgeAlyzer Solution:



- **Completely standalone** — all data stored and processed locally



- Accessible via customer network OR dedicated Supervisory Edge



- **Web browser interface** works on any device that can reach the IP



- **No external dependencies = no external failures**

Problem 3 – Proof of Concept is Difficult and Expensive

The Issue:



- POC requires significant customer IT resources before they even decide if they want it



- Security hurdles for systems on customer infrastructure



- Custom I/O configuration requires changes to customer equipment



- OEM equipment often requires vendor involvement to expose data

The EdgeAlyzer Solution:



- **Completely self-contained**










- Simple vibration POC: plug in → stick sensor on → system is running



- **Full functionality from day one** – no waiting for IT approvals

Novelty of the Approach

Innovation	What It Means
 True Edge Independence	No network required. No cloud dependency. Internet down? Still monitoring.
 Local Data Storage	Industrial-grade time-series database. Your data never leaves your facility.
 Extensive API & Integration	REST API, MQTT, Email – integrate when YOU choose to share data.
 Real AI Inference on Edge	Hailo-8 AI chip. 2-3ms inference. Neural networks, not thresholds.
 FiLM Conditioning	AI adapts to operating conditions. "High vibration at high RPM = normal."
 Self-Generating Training Data	Point-and-click model training. No data science expertise required.
 Universal Sensor Compatibility	0-10VDC, 4-20mA, 24VDC digital – any standard industrial sensor works.

The Competitive Landscape

Category	Competitors	Their Approach	EdgeAlyzer Difference
Cloud-Dependent Platforms	Augury, Senseye (Siemens), Uptake, Falkonry	Cloud processing. Internet 24/7. Subscription fees.	100% local. No internet. One-time purchase.
Traditional Vibration Vendors	SKF Enlight, Fluke 3563, Emerson AMS, Rockwell	Expensive. Complex software. IT infrastructure required.	Turnkey. Sensors + analytics + storage included.
Single-Sensor "AI" Products	Various IoT vendors	Claim "AI-powered" — actually just thresholds.	Multi-sensor correlation. True neural networks.
PLC/SCADA Add-ons	Native OEM monitoring	Limited to what OEM exposes. IT integration required.	Independent. Works alongside existing systems.

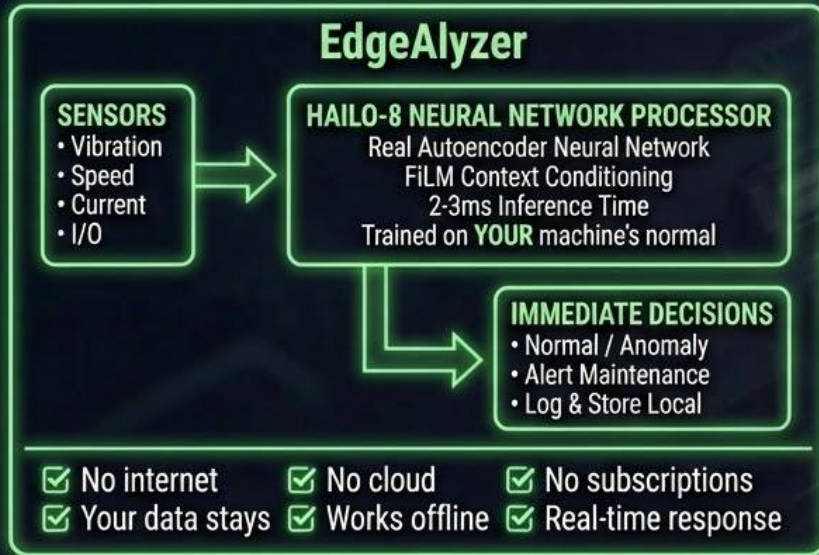
The AI That Lives On Your Machine – Not in the Cloud

Why EdgeAlyzer's Neural Network Changes Everything

The Industry Reality:

- ✗ Ships your data to the cloud – where their neural network runs on their servers
- ✗ Doesn't actually use AI – just thresholds and statistics rebranded as "machine learning"

EdgeAlyzer is Different:



What Makes This Possible:

Component	What It Does
Hallo-8 AI Accelerator	Purpose-built neural network chip – 26 TOPS dedicated to inference
Autoencoder Architecture	Learns compressed representation of "normal" – detects ANY deviation
FILM Conditioning	Modulates expectations based on speed/load – eliminates false alarms
Push-Button Training	Run your machine normally for 1-2 hours → model trains itself

The Result:

Metric	Cloud AI Competitors	EdgeAlyzer
Where AI Runs	Their servers	Your hardware
Inference Latency	1-60 seconds	2-3 milliseconds
Internet Required	Always	Never
Data Leaves Facility	Yes (all of it)	Never
Monthly Cost After Year 1	\$500-5000+/month	\$0
Works During Outage	No	Yes
Trained on YOUR Machine	Generic fleet models	Your specific equipment

Why This Matters:

Security: Your vibration signatures, operational patterns, and machine data are intellectual property. Cloud AI means sending this to external servers.

Reliability: When your internet goes down, cloud AI stops working. EdgeAlyzer keeps protecting your equipment.

Speed: Real anomaly detection happens in milliseconds. By the time cloud AI responds, damage may already be done.

Economics: No subscription fees. No per-sensor charges. No "you've exceeded your analysis quota" surprises.

Why Cloud Dependency is a Problem

Competitors Require

Constant internet connectivity

Data transmission to external servers

Monthly/yearly subscription fees

IT involvement (VPNs, firewalls, config)

Latency (seconds to minutes for results)

EdgeAlyzer Provides

No internet required

Data never leaves facility

One-time purchase

Zero IT burden

2-3 millisecond inference

Why Single-Sensor “AI” Doesn’t Work

The Technical Reality

A single data point cannot perform neural network inference.

True AI requires **multiple correlated measurements** to understand normal operation patterns.

Example

A current transformer alone cannot distinguish between:

- Normal high-load operation (high current = expected)
- Bearing failure (high current due to mechanical binding)

EdgeAlyzer Approach

- **Custom sensor packages** per application
- **Multiple correlated sensors** (vibration + temperature + speed + load)
- AI learns **relationships between sensors**
- Distinguishes “high vibration at high speed (normal)” from “high vibration at low speed (problem)”.

Why Operating Context Matters – FiLM Technology

The Problem with Generic AI

- Variable speed equipment: high vibration at high RPM is normal
- **Multi-mode machines:** rapid traverse vs. cutting have different profiles
- **Load conditions:** idle vs. loaded equipment behaves differently

Result: Constant false alarms →
Operators ignore the system →
Real problems get missed

EdgeAlyzer's FiLM Conditioning

- **Feature-wise Linear Modulation** adapts AI expectations based on operating state
- System understands: *“At 1800 RPM, this vibration level is normal”*
- **70-90% reduction in false positives** vs. context-unaware systems
- Validated in industrial applications (DMG MORI, Siemens)

Discoveries from Years of Field Experience

Discovery

- **Installation is the Real Problem**
- **Generic Solutions Don't Work**
- **Cloud Dependency Creates Fragility**
- **Operators Ignore False Alarms**
- **Data Ownership Matters**

Lesson Learned

- Technology works – IT hurdles kill projects. EdgeAlyzer eliminates barriers.
- "AI-powered" sensors that are just rebranded limit switches. We've seen them fail.
- IT updates, network changes, cloud outages take systems down. Edge independence is necessary.
- Expensive systems become background noise. Context-aware AI is the only solution.
- Customers locked into vendor platforms, unable to access own data. Local storage + open APIs.

Two Layers of Protection – ISO Standards + Neural Network AI

Industry Standards Meet Predictive Intelligence

Layer 1: ISO 10816 Vibration Severity Limits

Zone	Severitor	Lead Time	Icon	Action
Zone A	Good	< 0.11 in/s	✔	Green
Zone B	Acceptable	0.11-0.28 in/s	⚠	Yellow
Zone C	Unsatisfactory	0.28-0.71 in/s	⚠	Orange
Zone D	Unacceptable	> 0.71 in/s	⚠	Red

Built-in operator guidance:
One-click ISO presets

The Limitation: Limits Are Reactive



The Predictive Advantage



Layer 2: Neural Network AI – Earlier Detection

EdgeAlyzer's autoencoder learns YOUR machine's "normal"

ISO Limits	AI Pattern Detection
Severity thresholds	Subtle changes
After deterioration	At first deviation
Days/Hours lead time	Weeks/Months earlier

Real-World Example

Scenario	AI Neural Network Response	ISO
Bearing wear	Detects reconstruction error	No alert
Imbalance	Patterns diverging	No alert
Lubrication	Correlated shift	No alert

Why Both Layers Matter

ISO	AI
Industry-standard reference points	Machine-specific baseline learning
Regulatory compliance documentation	Early warning before limits are reached
Clear action thresholds	Continuous pattern monitoring
Operator confidence (known standards)	Subtle degradation detection

EdgeAlyzer combines the best of both: Operators understand ISO zones. The AI catches problems before you ever get there.

Smart Monitoring for Production Equipment — Process Inhibit

Built by Automation Engineers Who Understand Machines

The Problem with Continuous Monitoring on Production Equipment


Machine tools, CNC equipment, and production machinery have a fundamental challenge that data scientists rarely consider: **The machine vibrates differently when it's working vs. when it's idle.**

Machine State	Vibration Source	What It Tells You
Idle / Dry Run	Bearings, motors, spindles	Mechanical health
Active Machining	Cutting forces, material contact	Process — NOT mechanical health

The EdgeAlyzer Solution: Process Inhibit

EdgeAlyzer allows you to attach an analog signal (like a **spindle current transformer**) to each vibration sensor.

When the signal exceeds a threshold — indicating the machine is actively processing — AI monitoring automatically pauses.



Why This Changes Everything for Machine Tool Monitoring

If you monitor during machining:

- Training data becomes useless (mixing process noise with bearing signatures)
- AI triggers constant false anomalies (“cutting hard” looks like “bearing failing”)
- Operators disable the system → no protection at all

Traditional Approach	EdgeAlyzer with Process Inhibit
Train for a week of production	Train in 1-2 hours of dry run
Constant false anomalies during cuts	Zero false alarms during machining
Operators disable the system	System stays trusted and active
Can't distinguish process vs. failure	Clear separation of concerns

The Automation Engineer Advantage

This isn't something a data science team would think of. It comes from decades of experience standing next to machines, watching them work, and understanding that:

- A spindle cutting steel **vibrates 100x more** than an idle spindle
- That vibration is **completely normal** — it's supposed to do that
- The bearing failure you're looking for is **hidden under the process noise**

EdgeAlyzer was built by engineers with **100+ years of combined field experience** who understand the difference between “machine is working” and “machine is failing.”

Operator Visibility

The Group view shows when monitoring is paused and why:

Vibration Sensors ⓘ

Spindle Sensor: Z Velocity: 0.85 in/s X Velocity: 0.72	🔴 Process Inhibit (Spindle Current: 12.4A)
Axis Sensor: Z Velocity: 0.04 in/s X Velocity: 0.03	🟢 Monitoring Active

Operators know exactly what's happening — no mystery, no confusion, no ignored alerts.

Value Proposition – Maintenance & Plant Management

For Maintenance Teams

- **Predict failures** before catastrophic breakdown
- **Schedule maintenance** during planned shutdowns
- **Reduce false alarms** that waste investigation time
- **Simple interface** — no data science required

For Plant Management

- **Eliminate unplanned downtime** (can be thousand of dollars an hour)
- **One-time purchase** — no recurring cloud subscriptions
- **Fast ROI** — first predicted failure pays for the system

Value Proposition – IT & Corporate

For IT/Security

- **Zero network footprint** — isolated system
- **Data never leaves facility** — no cloud transmission
- **No STIG/security reviews** — not touching production systems
- **No servers to maintain** — self-contained platform

For Corporate/Finance

- **Predictable costs** — no subscription creep
- **Scalable** — from one machine to entire plants
- **Vendor independence** — open APIs, exportable data

Supervisory Edge – Scale to Full Plant

The Challenge:

How do you manage 20, 50, or 100+ monitored assets without drowning in complexity?

Traditional approaches require:

- Centralized servers with complex IT integration
- Cloud platforms with ongoing subscription costs
- Loss of edge independence (central failure = everything fails)

The Supervisory Edge Solution:

Communicates with all EdgeAlyzers over API and consolidates them into a **single unified interface** — while each EdgeAlyzer maintains complete standalone functionality.

Key Principle: If the Supervisory Edge goes offline, every EdgeAlyzer keeps monitoring independently. **No single point of failure.**

Supervisory Edge Capabilities

Feature	What It Does
Unified Operations Hub	Facility-wide dashboards, subgroup drill-downs, instant out-of-limit visibility
Secure Wi-Fi Hotspot	Bridges to EdgeAlyzers without customer enterprise network
Real-Time Monitoring	Continuously polls status, health, metrics, measurements, fault counts
Group & Facility Management	Organize by facility, department, production line — mirrors your plant organization
Built-In Security	Session-based auth, bearer-token access — authorized personnel only
Brandable Interface	Dynamic CSS for corporate branding without code changes
Offline Autonomy	Caches data, tolerates offline edges, degrades gracefully

From Single Machine to Entire Plant

Deployment Scale	Solution	Architecture
1-5 Machines	EdgeAlyzer standalone	Direct browser access to each unit
5-100 Machines	EdgeAlyzer + Supervisory Edge	Single interface, grouped by area
100+ Machines	Multiple Supervisory Edges	Regional hubs, each managing local EdgeAlyzers
Multi-Site	Supervisory Edge per site	Each site independent, optional corporate rollup

Start with one EdgeAlyzer, prove the value, scale to plant-wide coverage — without rearchitecting anything.

Why This Architecture Matters

vs. Alternative	The Problem	EdgeAlyzer + Supervisory Edge
Cloud-Centralized Systems	Cloud failure = entire monitoring down	Supervisory Edge failure = EdgeAlyzers keep running
Server-Based On-Premise	IT infrastructure, maintenance, backups	Same turnkey approach as EdgeAlyzer
Siloed Edge Devices	Multiple dashboards, no unified view	Single pane of glass

Best of both worlds: Centralized visibility with distributed resilience.

Technical Specifications

Specification	EdgeAlyzer Capability
Vibration Sensors	Up to 40 wireless, battery-powered, magnetic mount
Analog Inputs	Up to 40 channels (0-10V, 4-20mA)
Digital I/O	Up to 40 channels (24VDC)
AI Hardware	Hailo-8 accelerator (dedicated neural network chip)
Inference Speed	2-3 milliseconds per sample
Data Storage	Local database
Interface	Web-based (any browser)
Connectivity	Ethernet, optional Wi-Fi, isolated network
Power	Standard 120VAC outlet
Enclosure	18"×12" industrial-grade
Integration	REST API, MQTT, Email alerts, Digital outputs

Competitive Comparison Summary

Feature	EdgeAlyzer	Cloud AI Platforms	Traditional Vendors	Single-Sensor "AI"
True Neural Network AI	✓ Autoencoder + FiLM	⚠ Varies (unclear)	✗ Thresholds	✗ Marketing only
Edge Processing	✓ 100% local	✗ Cloud required	⚠ Mixed	✗ No processing
Internet Required	✗ No	✓ Yes (always)	⚠ Often	✗ N/A
Subscription Fees	✗ One-time	✓ \$\$\$\$ / year	⚠ Often	⚠ Sometimes
IT Integration	✗ No	✓ Yes (extensive)	✓ Yes	✗ Minimal
False Alarm Reduction	✓ 70-90% lower	⚠ Varies	⚠ Poor	✗ Very poor
Context Aware	✓ FiLM conditioning	✗ Usually not	✗ No	✗ No
Scalable to Plant	✓ Supervisory Edge	✓ Cloud native	⚠ Complex	✗ Single machine
Distributed Resilience	✓ Each edge independent	✗ Single point failure	⚠ Varies	✓ Isolated

Questions to Ask Other Vendors

Question	What You're Looking For
Does it require internet connectivity?	If yes → not truly edge-based. What about network outages?
Can it handle variable-speed equipment without false alarms?	If no → doesn't understand operating context.
Is the sensor package customized for my application?	If no → generic solution may miss YOUR failure modes.
Can you show me the neural network architecture?	If they can't → might just be thresholds with "AI" marketing.
What are the ongoing costs after year one?	Cloud subscriptions add up. EdgeAlyzer is one-time.
Who owns the historical data? Can I export it?	If unclear → may be locked into their platform forever.
If your central system goes down, do I lose all monitoring?	If yes → single point of failure.
EdgeAlyzer's answers: ✓ No, ✓ Yes, ✓ Yes, ✓ Yes (documented), ✓ None, ✓ You/Yes, ✓ No	

The Bottom Line

***EdgeAlyzer is not just another 'AI-powered' sensor.*"**

It's a **complete predictive maintenance system** built by industrial automation experts who understand:

- **✓** How equipment actually fails in the real world
- **✓** What sensors are needed for accurate prediction
- **✓** How to implement true AI on the edge
- **✓** How to reduce false alarms through operating context awareness

Call to Action

No Cloud. No Subscriptions. No IT Headaches. No False Alarms.

From a single machine to an entire plant — with the **Supervisory Edge**, scale your predictive maintenance program while maintaining the independence and resilience that makes EdgeAlyzer unique.

Just reliable, industrial-grade predictive maintenance that works.

Ready to eliminate unplanned downtime with true edge AI?

Contact FGR Automation
to discuss your specific application