

INDUSTRY · DIAGNOSTIC DELIVERABLE

Industry Assessment · Manufacturing Vertical — SAMPLE

4-week vertical assessment · AI & automation use cases for manufacturing

Client	Empresa Anônima Indústria S.A.
Engagement ID	#ANV-IND-2604-0312
Period	Q2 2026 · April 06 – May 04, 2026
Issued by	Anuvia · Mila Vernazza, Principal Consultant
Classification	SAMPLE · Confidential · Not for distribution

This document is a sample preview of an Anuvia diagnostic deliverable. All figures, customer names, and findings are illustrative and anonymized. No real-client data is contained within.

anuvia.com.br · contato@anuvia.com.br

VERTICAL CONTEXT

Manufacturing • Brazil • mid-market

The Brazilian mid-market manufacturer (R\$ 200M-R\$ 1B revenue band) operates within a specific set of constraints: a maturing Industry 4.0 adoption curve, mixed OT/IT estate with legacy PLCs, increasing pressure on energy efficiency, and the LGPD plus sector-specific regulations (NR-12 for machine safety, ABNT operational standards).

Regulatory environment

Regulation	Scope	Impact on AI / automation
LGPD	Personal data of employees, suppliers, customers	Employee operational data; supplier data lawful basis
NR-12 (MTE)	Machine & equipment safety	AI-driven control loops require fail-safe certification; advisory-only mod
ANVISA / sector-specific	Where applicable (food, pharma)	Process-change traceability mandatory.
MROSC / fiscal reqs	Documentation, audit trails	Any AI-driven decision affecting fiscal/operational reporting must log in

Market patterns observed

- OEE (Overall Equipment Effectiveness) avg in segment: 58-68% — significant headroom.
- Predictive-maintenance adoption is partial: 38% of mid-market manufacturers run sensors but only 12% drive a closed loop with CMMS.
- Energy cost share growing — 7-12% of cost-of-goods, up from 5-7% three years ago.
- Supply chain volatility persists — forecast accuracy below 65% for most.

USE CASE RANKING

Top 5 AI / automation use cases • manufacturing

Rank	Use case	Impact horizon	Capex	Verdict
#1	Predictive maintenance • sensors + anomaly + CMMS	12-18 mo	Med	Build • phased
#2	Quality vision (defect detection at the line)	6-12 mo	Med	Build • pilot first
#3	Energy optimization • HVAC + compressors	9-18 mo	Low	Buy + tune
#4	Demand forecasting refresh	6-12 mo	Low	Buy + tune
#5	Document AI • purchase orders + invoices	3-6 mo	Low	Build • light

Scoring rationale

- Predictive maintenance ranked #1: highest expected NPV across the band, strong data foundations from existing PLC instrumentation, regulatory path well-trodden.
- Quality vision ranked #2: best fit for incremental rollout (line-by-line); strong evidence of ROI in adjacent verticals.
- Energy optimization ranked #3: lowest capex, but ROI sensitive to local energy tariff schedule.
- Demand forecasting ranked #4: well-defined; many mature vendors; benefits diminish as upstream data quality plateaus.
- Document AI ranked #5: smallest ticket but quickest payback — "first win" candidate.

USE CASE DEEP-DIVE · #1

Predictive maintenance · sensors + anomaly + CMMS

A closed-loop predictive maintenance program reduces unplanned downtime by 18-32% in comparable mid-market manufacturing settings (CARS-D consortium benchmarks, 2024). The system requires three components stitched together — sensors generating high-frequency vibration / temperature / current data, an anomaly-detection layer, and integration with the existing CMMS to convert anomalies into prioritized work orders.

Component stack

Layer	Recommended	Capex / vendor
Edge sensors (vibration + temp)	Wireless IIoT sensors · 200 units phase 1	R\$ 380k (one-time)
Edge gateway	Industrial gateway + buffering · 8 units	R\$ 140k
Data ingest	AWS IoT Core · sa-east-1	Opex: R\$ 4k/mo
Anomaly model	Isolation forest baseline + per-asset tuning	Build (internal)
CMMS integration	REST adapter to existing CMMS (Engeman / MV2)	Build (4 wks)
Visualization	Grafana dashboards + ops mobile alert	Build (2 wks)

Risk register

- Sensor mount quality is a top risk — 15% of pilot deployments report data drift from poor mounting.
- False-positive alerts erode operator trust within 90 days if not tuned.
- Spare-parts logistics — CMMS-triggered work orders only deliver ROI if parts are stocked.

ROI MODEL · USE CASE #1

Predictive maintenance — 36-month model

Assumptions (explicit)

Assumption	Value	Source
Current unplanned downtime	412 hr/yr	Client CMMS export 2024-2025
Downtime cost per hour	R\$ 28,000	Client finance · GM-throughput-based
Expected downtime reduction (steady-state)	24% (mid of 18-32% band)	CARS-D 2024 benchmark
Time-to-steady-state	Month 9	Anuvia engagement pattern
Sensor + gateway capex	R\$ 520k	Vendor quote · BR market
Annual opex (cloud + software)	R\$ 96k	AWS IoT Core + CMMS integration support
Spare parts buffer increase	R\$ 80k one-time	Stock-level adjustment

36-month cash flow (R\$ thousands)

Year	Capex	Opex	Downtime savings	Net
Y1	-600	-96	+1,389 (partial)	+693
Y2	0	-96	+2,769	+2,673
Y3	0	-96	+2,769	+2,673
Total	-600	-288	+6,927	+6,039

Payback: **Month 9**. NPV (12% discount, 36mo): **R\$ 4.7M**. IRR: **183%**.

COMPLIANCE CONSIDERATIONS**LGPD + sector regulations****LGPD**

- Operational sensor data is generally non-personal — outside LGPD primary scope.
- Operator identification (badge → machine → output) is personal data — requires lawful basis (legitimate interest or contract).
- Build a data inventory + DPIA for use cases that link operator identity with performance metrics.
- Retention policy: aggregate operational data may be retained indefinitely; operator-attributable data ≤ 24 months unless compelling justification.

NR-12 (machine safety)


- AI driving direct machine actuation must integrate with NR-12-compliant fail-safe interlocks.
- Recommend advisory-only deployment for the first 12 months — operator-in-loop.
- Document the change-management process per ABNT NBR ISO/IEC 27001-aligned controls.

Audit-trail requirements

- Every AI-driven recommendation logged with: input features (hashed), model version, output, downstream action taken, operator confirming.
- Retention 5 years (alignment with fiscal/operational audit standards).
- Quarterly model-review minutes archived.

IMPLEMENTATION ROADMAP

12-month phased rollout

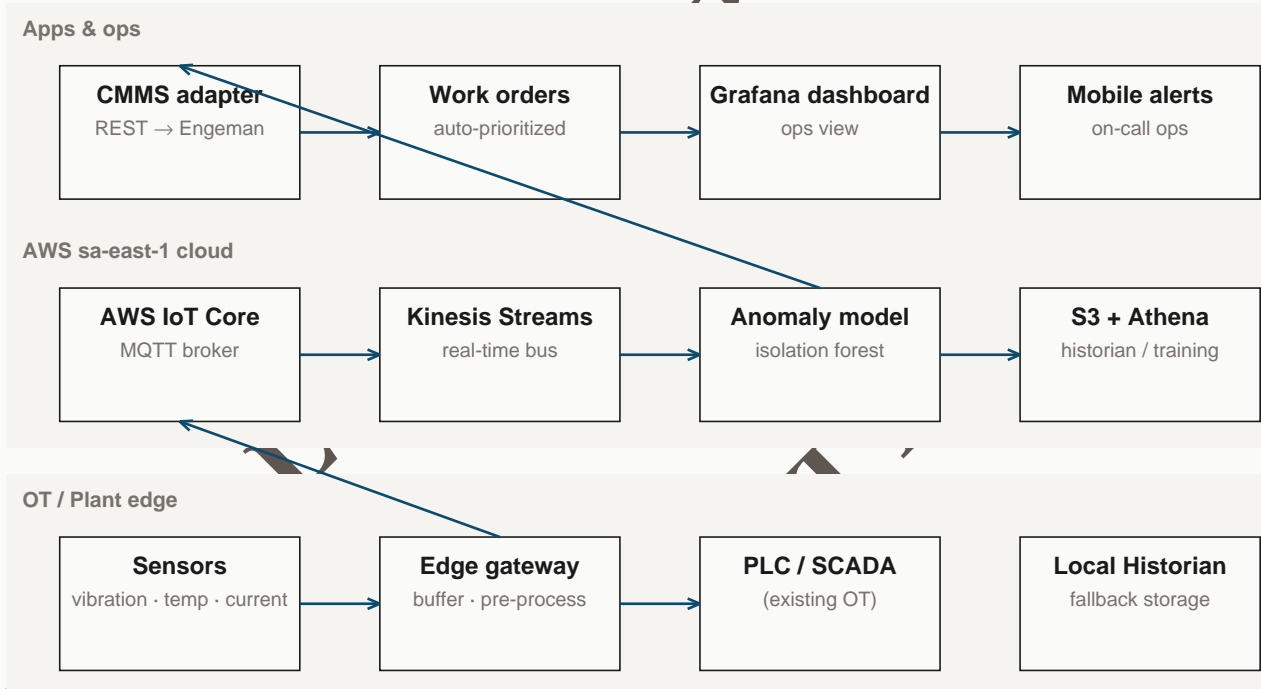


Gate criteria

- Pilot → Rollout: defined success metric achieved on 1 line for 6 weeks; runbook complete.
- Rollout → Scale: zero P0 events for 4 weeks; ops team owns operations.
- Scale → Operate: financial review against business case; quarterly committee.

REFERENCE ARCHITECTURE

Predictive maintenance • target topology



Topology notes

- All control-plane traffic stays in sa-east-1 (LGPD + data sovereignty).
- Edge gateways buffer up to 24h to survive WAN outages.
- Anomaly model runs serverless (Lambda + SageMaker endpoint for batch retraining).
- CMMS integration is one-way (advisory) for the first 12 months.
- Mobile alerts use existing on-call rotation; no new tooling.

— End of report — anuvia · contato@anuvia.com.br · anuvia.com.br