FREE GUIDE · TECHFLOW INDUSTRIAL

Industrial Data Platform in 7 Steps

From raw sensor data to actionable information — no vendor lock-in, no €40,000/year license costs.

> *"Byte 0, Value 50 — what does that mean to a data lake? Nothing, without context."*

AVEVA CONNECT	IDP OPEN SOURCE STACK
€40.000	**€8**
per year · vendor lock-in	per month · fully open source

This guide teaches you HOW — the 7 steps to build your own Industrial Data Platform. We deliver the implementation.

Johannes Sanderse
OT Engineer · 15 years in factories (Siemens PLC/SCADA, historians, DCS)

techflow24.com

Your factory doesn't have a data problem.

It has a context problem.

Every PLC publishes data. But data without context is noise. What your historian receives every second looks like this:

Raw sensor data — what your system receives:

```
TIC-001 = 67.3
PIC-001 = 4.47
FIC-001 = 248.1
BatchCounter = 142
← What does this mean? Nobody knows. This is the problem.
```

The value **67.3** is useless without knowing: from which machine? In which engineering unit? What is normal? What is an alarm? **Context is what AVEVA is trying to sell you for €40,000 per year.**

Data → Information ladder

DATA	`TIC-001 = 67.3`
	↓ Step 3: MODEL (add ISA-95 context)
INFORMATION	`WC-PAST-01 · Temperature · 67.3°C · Normal: 55–75°C`
	↓ Steps 4–5: STORE + ORCHESTRATE
INSIGHT	`Temperature rising towards alarm · trend +2°C/hr`
	↓ Steps 6–7: VISUALIZE + DISTRIBUTE
ACTION	`Operator adjusts cooling · N8N sends Teams notification`

The 7 steps in this guide walk you through every rung of the ladder — from raw tag to actionable information.

The 7-step pipeline

From OPC-UA sensor data to actionable information — all on a €8/month VPS.

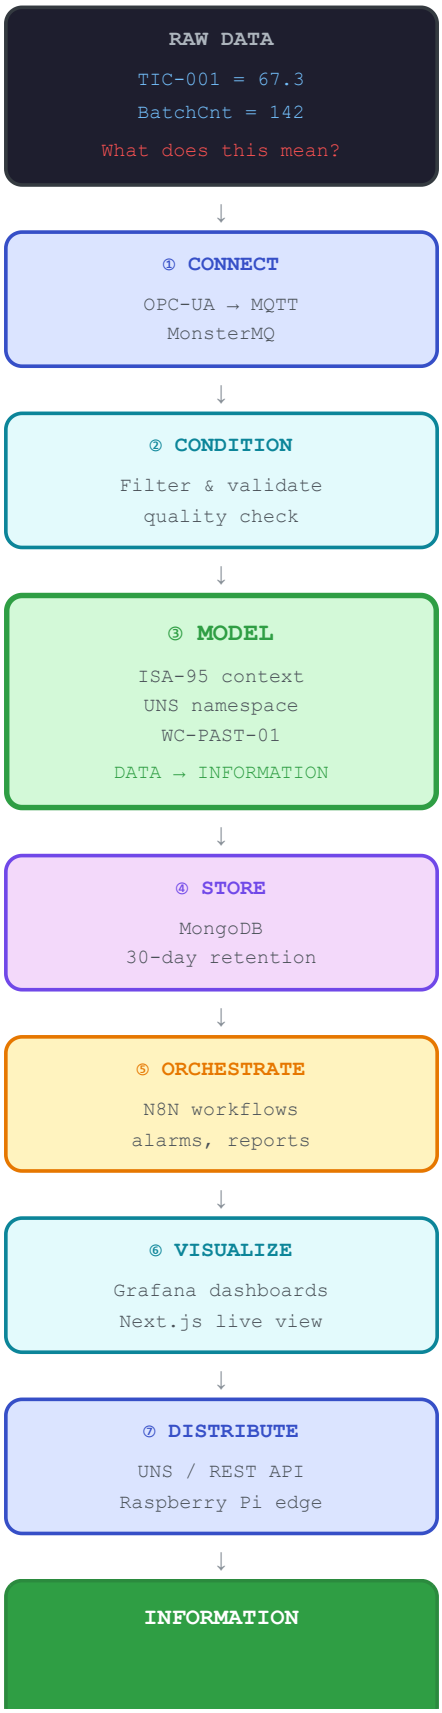

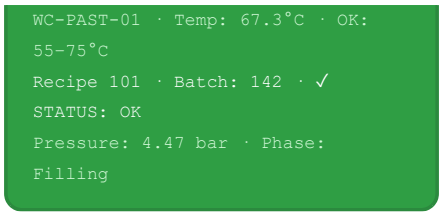

"Byte 0, Value 50 — what does that mean to a data lake?
Nothing, without context."

Open source · €8/month VPS · Zero license costs

Figure 1: The 7-step IDP pipeline — from raw sensor data to actionable information. Step ③ MODEL is the critical point.

The 7 Steps

1

STEP 1

CONNECT — Ingesting data

Connect to the data source of the PLC or SCADA. OPC-UA is the industry standard: every modern system speaks it. MonsterMQ acts as OPC-UA client and MQTT broker in a single container — no separate bridge needed.

OPC-UA (asyncua) | MonsterMQ | MQTT | Modbus (future)

2

STEP 2

CONDITION — Making data reliable

Remove duplicate values, filter noise, convert engineering units. A PLC publishes once per second — but not every value changes. Storing only changed values saves 60–80% of storage space.

MonsterMQ flow engine | N8N validation

3

STEP 3 — CRITICAL POINT

MODEL — Adding context (DATA → INFORMATION)

This is where data becomes information. Use the ISA-95 hierarchy (Enterprise → Site → Area → WorkCenter → Tag) to contextualise every sensor value. A JSON model describes every machine: what it does, what is normal, which alarms exist, and how it relates to other machines.

The four namespace types: *Definitional* (what is it?), *Functional* (what does it do?), *Informative* (what does it measure?), *Ad-hoc* (temporary/batch).

ISA-95 JSON model | UNS (Unified Namespace) | MQTT topic structure

4

STEP 4

STORE — Storing information

MongoDB stores all MQTT messages in the `plc_data` collection with automatic timestamps. Retention 30 days. No InfluxDB needed — MongoDB handles mixed data (sensor values, batch records, alarms, shift logs) in a single collection.

MongoDB | MonsterMQ archiving | 30-day retention

5

STEP 5

ORCHESTRATE — Making information flow

N8N automates factory operations: starting production orders, sending alarm notifications to Teams, processing barcode scans, emailing shift reports. The IDP stack defines 6 workflows for a complete packaging line.

N8N workflows | MonsterMQ webhooks | Teams / email

6

STEP 6

VISUALIZE — Making insights visible

Grafana for time-series analysis and ad-hoc queries (via FastAPI JSON datasource). Next.js for live dashboards that refresh every 5 seconds — ideal for operators on the shop floor. Grafana template variables: one dashboard for 50 machines.

Grafana | Next.js | FastAPI

7

STEP 7

DISTRIBUTE — Enabling action

The Unified Namespace (UNS) makes data available to every system that needs it — ERP, MES, cloud, edge. MonsterMQ on a Raspberry Pi forwards local data to the VPS. REST API for external integrations. MQTT retain flag for startup states.

UNS / MQTT | REST API | Raspberry Pi edge

The technical stack

TechFlow Plant A — components

OPC-UA Simulator (3 PLCs)
MonsterMQ (MQTT broker + OPC-UA client)
MongoDB (plc_data, 30-day retention)
N8N (workflows: alarms, orders, reports)
FastAPI + Grafana + Next.js dashboard
Traefik (reverse proxy + SSL)

Simulated factory

PLC_01: Pasteurisation (temp, pressure, flow)
PLC_02: Conveyor & labeller (3 motors)
PLC_03: Batch control (recipe, phase, counter)
N8N: 6 factory workflows
ISA-95 JSON model (complete)
Fully open source · €8/month

Stack comparison

SOLUTION	COST/YEAR	OPC-UA	OPEN SOURCE	LEARNABLE
AVEVA Connect	€40.000	✓	No	Vendor training
Commercial MQTT stack	€6.500–€20.000	✓	No	Vendor training
IDP Open Source Stack	**€96/year (VPS)**	✓	100%	HOW — live demo

N8N Workflows in the stack

WF-001
Start Production Order
Webhook → MongoDB → response

WF-002
Alarm Handler
MonsterMQ → MongoDB → Teams

WF-003
Material Scan
Barcode → lookup → inventory log

WF-004
Product Scan
End-of-line → batch counter

WF-005
Shift Report
Cron 3x/day → KPIs → email

WF-006
Fault Detection
FaultBits → decode → Teams

Total cost open source stack: €8/month VPS · Zero extra license costs.
Compare: AVEVA Connect €40,000/year · Commercial MQTT stack €6,500–€20,000/year.

Ready to build it yourself?

Join the community of PLC/SCADA engineers building their own Industrial Data Platform — step by step, open source, no vendor lock-in.

1	2	3
Join the community	Sign up for live workshop	Fork the GitHub repo & start

techflow24.com

GitHub: github.com/cftservices/Industrial_Data_Platform