

CUSTOMER SUCCESS STORY

The factory floor got smarter. The data never left it.

Predictive maintenance, inline quality inspection, a plant-scale digital twin, and robotics — four programmes across a global manufacturing network, all on **UPC AI Factory**, with process IP that never leaves.

52–58%

Manufacturing AI adoption now sits in the 52–58% range, with spend growing fastest in predictive maintenance, quality control, and physical AI. The gap is execution: about two-thirds of maintenance teams plan AI adoption, but only a third have implemented it.

BUILT ON UPC AI FACTORY

THE SITUATION

Pilots on one line, stuck before the second

Each programme proved itself on one line: vibration models caught bearing failures early, vision models caught defects the eye missed. Scaling to the network is where they stalled:

- **Process IP could not leave.** Defect imagery, recipes, and tolerances are crown-jewel IP, and legal would not put them on shared infrastructure.
- **Inline inspection runs at line speed.** A decision made after the part has moved is worthless; a round trip to a distant region was never viable.
- **Sensor and video volumes were enormous.** Moving telemetry out of plants generated egress charges that grew with every line added.
- **Simulation demand was spiky.** Digital-twin runs needed large GPU bursts, making on-demand pricing impossible to plan around.

THE DECISION

Put the factory's AI where the factory is

The network standardised on one private AI platform — operated by UnitedLayer, sited near the plants, governed by the manufacturer.

“Our defect images and process recipes are the business. Once we stopped exporting them to run a model, scaling to the whole network became an engineering problem, not a legal one.”

— VP, MANUFACTURING TECHNOLOGY

CUSTOMER PROFILE

INDUSTRY

Global industrial manufacturer, multi-plant discrete

AI WORKLOADS

Predictive maintenance · Quality inspection · Digital twin · Robotics

FOOTPRINT

Two regional data centres plus a DR site

TENANCY

Isolated vDCs — Operations, Quality, Simulation

TIME TO PRODUCTION

12 weeks, end to end

GOVERNANCE

Process IP in-boundary, OT/IT segregation

SOLUTION

UPC AI Factory, fully managed

WHAT WAS DEPLOYED

One platform, from the sensor to the simulation

AI services

OpenAI-compatible Router, Model Registry for per-line versions, managed Vector DB for maintenance history, and GPU Workbench for simulation engineers.

Compute

Dedicated GPU pools with MIG partitioning, keeping bursty digital-twin simulation away from the steady inference that inspection and robotics depend on.

Data

High-performance storage for sensor telemetry, defect imagery, and twin state, co-located with GPU so plant data is processed where it lands.

Governance & operations

OT/IT segregation, backup, DR, and a cyber vault for process models; observability and AIOps into existing tooling, plus FinOps dashboards by plant.

THE RESULTS · YEAR ONE

↓ **42%**

Lower TCO for AI workloads

Egress and cross-region movement designed out

↑ **40%**

Faster model training

Single-tenant, RDMA-native GPU fabric (35–40%)

<100ms

p95 inference latency

Inspection decisions that keep pace with the line

100%

Process IP and models inside the tenant boundary

Defect imagery, recipes, and tolerances never left — removing the legal blocker that had held back network rollout.

~**60%**

Less platform toil

One operated stack replaced per-plant tooling, so a model proven on one line deploys to the next without a rebuild.

Fixed

Cost per plant and per programme

Maintenance, quality, simulation, and robotics each carry a known cost, making line-level business cases straightforward.



Spotlight · The digital twin

The twin was the programme that made the others worth scaling — and the one that could not have run on metered infrastructure.

Simulation at plant scale

Fixed GPU capacity turned twin fidelity into an engineering decision instead of a budget request, so line changes are modelled before they are built.

A closed loop with the line

Live telemetry keeps the twin current, and maintenance and quality models are validated against it before they reach production.

The IP stays inside

A twin encodes how the factory actually works. Keeping it in-boundary meant the most sensitive asset never became an external dependency.

WHY THIS MATTERS IN MANUFACTURING

Process IP kept in-boundary

OT/IT segregation

Line-speed inference

No egress on sensor & video

Per-plant cost visibility

Repeatable line-to-line rollout

BOTTOM LINE

UPC AI Factory took AI from a single proven line to the whole network — at line speed, inside the boundary, on a cost base the plant can plan.

Ready for enterprise AI without compromise?

Book a walkthrough, or ask us to scope a private AI factory for your workloads.

UPC AI FACTORY



EMAIL

contact@unitedlayer.com



PHONE

+1-888-853-7733

[UNITEDLAYER.COM](https://unitedlayer.com)