

SECTOR ANALYSIS • MANUFACTURING & INDUSTRIALS

The Line Either Runs Or It Does Not

Ask a plant manager what a stopped line costs per minute and they will answer from memory. Ask them what a degraded one costs and they will look puzzled — because there is no such thing.

THALAMUS ADVISORY • ENTERPRISE RESILIENCE PRACTICE

01 — THE SIGNAL

A number everyone already knows

Manufacturing is unusual in this series: the business case needs no construction. The cost of downtime is a figure the plant manager can quote, the CFO has already modelled, and nobody argues about. What is missing is not justification — it is the ability to act on it in software terms.

US manufacturing output runs to roughly \$2.9 trillion, about 10% of GDP. The estate is bimodal: MES and SCADA on the plant floor operating on one set of assumptions, ERP and PLM in the enterprise operating on another, and a widening industrial IoT layer between them feeding predictive maintenance models.

02 — THE MECHANISM

Recovery is physical

Software incidents in manufacturing have a property they do not have elsewhere: restarting does not fix anything.

A paint line, a furnace or a batch reactor cannot be restored by restarting a process. Work in progress may be scrapped. Restart sequences take hours and have their own failure modes. The gap between “the software is healthy again” and “the line is producing again” is where the real cost sits, and almost no monitoring measures it.

There is no partial degradation on a production line. There is running, and there is a number the plant manager can quote from memory.

The second pattern is the reliability-culture split. Traditional manufacturers are becoming software operators — connected vehicles, telematics, fleet services, customer platforms — and these have consumer-software reliability expectations inside companies whose entire operational culture was built around plants. The two have different tempos, different definitions of urgent, and usually one shared IT function trying to serve both.

03 — THE COST

Where it actually lands

SURFACE	HOW IT PRESENTS
Line stoppage	Directly quantified per minute, with restart time and scrapped work in progress on top of the stoppage itself.
Seasonal and launch peaks	Agricultural equipment has planting and harvest windows measured in days. A platform failure during harvest cannot be compensated later.
Traceability gaps	Aerospace and regulated manufacturing require provable chain of custody. A silent data-loss incident becomes a compliance event discovered in an audit months afterwards.
OTA and fleet risk	Vehicle software shipped over the air is a recall-class event when wrong, not a rollback.
Dealer and partner SLAs	Parts and dealer systems carry contractual uptime commitments to independent businesses whose own revenue depends on them.

04 — THE AI STAKES

Velocity meets a physical blast radius

Manufacturers are adopting AI across quality inspection, predictive maintenance, supply chain planning and increasingly vehicle autonomy. Each places a probabilistic component somewhere with a physical consequence.

The reliability risk is deploy velocity. Software-defined vehicles and connected products push change rates that plant-floor culture has never had to absorb, and at high deploy velocity *“what changed just before onset”* is almost always the answer — but only if the change correlation is anchored to when the condition began rather than when someone started investigating.

Industrial estates are also a standing ransomware target, and the sector’s tolerance for operational interruption is the lowest in the economy, which is precisely why it is targeted. An estate that cannot distinguish a silent partial failure from an intrusion is not in a position to respond to either.

05 — THE REMEDY

Close the gaps before the window

- **Audit for missing resilience patterns in the enterprise layer.** Absent circuit breakers, unbounded retries, shared pools and non-idempotent writes on retried paths are findable in source, at leisure, long before a peak.
- **Derive configuration from observed traffic** rather than from defaults set during commissioning and never revisited.
- **Separate objectives by domain.** Plant systems and connected-product services should hold different SLOs with independent error budgets rather than one standard that is too slow for one and too loose for the other.
- **Detect silent partial failure.** Baseline-relative regression on event classes, and reconciliation counts between pipeline stages, so a traceability gap surfaces as an incident rather than an audit finding.
- **Guardrail anything touching physical systems.** Automated remediation against MES, SCADA or OTA paths should escalate by policy, not by convention.

06 — IN PRACTICE

Hosted, or customized for your estate

Hosted on Thalamus AI Cloud

Well suited to the enterprise and customer-facing half — dealer systems, parts and aftermarket, telematics analytics, supply chain portals and connected-product platforms. Managed instances remove the infrastructure lead time, and dedicated tenancy is available where supplier or customer data cannot share a store.

Customized for your enterprise

Plant-floor observability is self-hosted in practice, and usually on-premise per site or per region. The platform runs entirely inside your network with no egress. Customization for manufacturers generally covers: collectors deployed per plant with aggregation upward, so a site keeps working when the WAN does not; SLO templates per asset and line class rather than per microservice; guardrail policy that makes anything touching MES or SCADA escalate to a named human; and gap tickets routed into your existing maintenance and change windows rather than competing with them.




FIGURE 1 — THALAMUS SRE · GAP TICKETS

Resilience gaps raised as scheduled work. Structural findings — *liveness/readiness probes missing or shallow, shared thread/connection pool across dependencies, non-idempotent write path on a retried operation* — alongside traffic-derived tuning recommendations, each priced by priority and routed to ServiceNow or Azure DevOps. For a manufacturer this is the mechanism that converts reliability from an incident response activity into planned maintenance.




FIGURE 2 — THALAMUSTRACE • GOLDEN SIGNALS

Per-service health, not an estate average. Failure rate, p50, p95, throughput and Apdex for every service with a letter-graded health score. One service at 9.63% failure and grade D against a fleet average of 93 is the pattern that matters in a multi-site estate — a single plant or line degrading while the aggregate stays comfortably green.

07 — GETTING THERE

How Thalamus Advisory helps

ENGAGEMENT

WHAT IT PRODUCES

Stoppage Attribution

3–4 weeks

Your last 12 months of unplanned stoppages classified by whether the root cause was software, and what the software-attributable minutes actually cost. Usually the first time anyone has separated the two.

Pre-Peak Resilience Audit

4–6 weeks

Repository and configuration audit ahead of a seasonal or launch window, raised as scheduled work inside your existing change calendar.

Traceability Assurance

4 weeks

Reconciliation and baseline-relative detection across serialization and chain-of-custody pipelines, so a silent gap is an alert rather than an audit finding.

Dual-Domain SLO Design

4–6 weeks

Separate objectives and error budgets for plant systems and connected-product services, under one investigation practice.

THE COUNTER-ARGUMENT, MADE HONESTLY

Most plant-floor systems are not where this starts, and a vendor proposing to instrument your MES in month one should be declined. Change control exists for reasons that predate software, and the validation cost is real.

The sequence that works is enterprise layer first — dealer, parts, telematics, supply chain — where the change cost is low and the value is provable. The plant-floor conversation should happen later, with evidence, and probably not with observability as the opening move.

ONE NUMBER

Zero. Minutes of acceptable partial degradation on a production line.

Every other sector in this series has a degraded mode it can operate in while it investigates. Manufacturing does not, which is why the economics favour prevention so heavily over response.

Where to start on Monday

Take your last ten unplanned stoppages. For each, ask whether the root cause was **mechanical, process, or software** — and if software, how long the diagnosis took relative to the restart.

Most manufacturers have never separated those three. The software share is usually larger than expected, and it is the only one of the three that responds to this kind of investment.

Thalamus Advisory

Enterprise resilience for the age of autonomous systems. We help organisations build operational practice that stays trustworthy when the systems making decisions are no longer entirely human.

THE RESILIENCE NEWSLETTER · ISSUE 08 · SEPTEMBER 2026

Every figure we print carries its assumptions. Where we model, we say so.

Written for engineering leaders. Forward it to one.