



THE INDUSTRY EDITION · ISSUE 01 · 18 SEPTEMBER 2026

Smarter production. Smaller attack surface.

A practical guide to AI-assisted quality and maintenance —
without giving software the authority to run your plant.

MANUFACTURING & FOOD AND BEVERAGE

THE OPPORTUNITY

One workflow, not a factory-wide transformation.

THE CYBER QUESTION

What can the assistant reach beyond approved data?

THE ACTION

A 30-day pilot and a business case you can challenge.

In this issue

Seven parts, written to be read in order at a production meeting. Research cut-off: 11 September 2026.

-
- 01 The briefing — the useful factory AI is not always a robot.**
 -
 - 02 AI in practice — one line, one question, one accountable owner.**
 -
 - 03 Cyber front — secure the connection, not just the conversation.**
 -
 - 04 The 30-day playbook — prove value before granting authority.**
 -
 - 05 The decision desk — count recovered capacity, not imaginary savings.**
 -
 - 06 Monday's agenda — leave the meeting with a decision, not a demo.**
 -
 - 07 Sources & editorial notes — evidence you can inspect.**
-

HOW TO READ THIS EDITION

Reported examples are attributed to their sources. The SME scenarios, pilot plans and decision questions are original editorial proposals, not verified Agentsstack customer results. All cost examples use explicit assumptions, not market prices or promised returns.

Signals are verified to 11 September 2026. They are a dated briefing, not claims about news that will occur during publication week.

The useful factory AI is not always a robot.

For owners, plant managers and technology leaders: choose a narrow operational problem and make the control boundary part of the investment.

The editor's view

A good first AI project should be explainable at the morning production meeting. Which recurring task is too slow? What information does the team already trust? Who remains responsible when a recommendation is wrong? Those questions are more useful than starting with a shopping list of models, cameras and robots.

Start where people already decide

For a smaller food or beverage producer, our recommended starting points are maintenance-document search, packaging-

image triage or a draft explanation of yesterday's downtime. Keep product release, recipe changes and machine actuation outside the first pilot. The aim is useful assistance with a reversible failure mode, not unattended production.

The test for this edition

Select one process with a named owner and a measurable baseline. Then ask whether the same improvement is possible with clearer procedures, a basic dashboard or rules-based automation. AI needs to beat that simpler alternative, not just the current frustration.

O1 / A production example, not an SME promise

Computer Weekly's August 2026 public synopsis describes substantial robotic assembly work at Hyundai's Singapore factory. It illustrates physical-world AI at industrial scale; it does not establish the cost or return for a smaller bottling line. [1]

O2 / The controls around the model matter

ASD's 11 September guidance examines the software layer connecting an AI agent to

tools and systems. Its warning is practical: risk depends on what the agent can access and do, not just on how its model answers.

[2]

O3 / Critical-infrastructure guidance is evolving

NIST is developing a trustworthy-AI profile for critical infrastructure. The project includes a concept note and stakeholder engagement; it is not a finished mandatory rule. Treat it as a direction of travel for risk discussions. [3]

One line. One question. One accountable owner.

The best first use case is the one your team can evaluate before it can change anything.

Begin with the decision, not the model

Consider a packaging line where technicians repeatedly look up approved setup and maintenance instructions. A retrieval assistant could return relevant passages from controlled documents, with document version and location. The technician still decides what to do. This is our proposed SME use case, not a reported Agentsstack client deployment.

Make the evidence visible

Give every answer a short evidence panel: the source, its approval date, the equipment

variant and any conflicting instructions.

Require the assistant to stop when a document is missing or the machine identifier does not match. A fluent answer without the correct asset context should be treated as a failed answer.

Keep the alternative honest

Compare the assistant with a well-indexed document library. If ordinary search already gets the technician to the right procedure quickly, adding a generative answer may introduce review work rather than remove it. Test the extra value on real tasks, not on a scripted demonstration.

Where visual inspection can fit

A second candidate is packaging-image triage: highlighting a suspected label mismatch or visible defect for a quality specialist. Start with historical images and then observation-only operation. This proposed pilot does not replace established inspection, sampling, food-safety checks or release authority.

Design for mistakes you care about

Separate missed defects from false alarms. A system that flags every image can appear cautious while overwhelming the reviewer; a

system that rarely flags anything can hide a serious miss. Record both, along with lighting, product variant and image quality. Include an explicit “unable to assess” outcome.

Why the factory context matters

NIST’s OT guide emphasizes that systems interacting with physical processes have distinctive safety, reliability and performance needs. [4] Our implication: the initial assistant should read an approved copy of information, not hold a path to a controller, interlock or recipe database.

The first deliverable is not autonomy. It is evidence that assistance improves a real decision.

AGENTSSTACK PERSPECTIVE

Secure the connection, not just the conversation.

A well-behaved model is not a substitute for a deliberately constrained system.

An illustrative failure

Imagine a maintenance assistant reading a supplier PDF. The document contains instructions directed at the assistant rather than the technician. The risk is not merely an odd reply; it is whether the assistant can use a connected tool to retrieve unrelated material or alter a system. This is a hypothetical scenario, not an incident report.

Treat documents as data

OWASP identifies indirect prompt injection through external content and warns that retrieval does not eliminate this risk. [5] For

the proposed plant pilot, supplier files enter a review area before indexing. A document's contents must never define the assistant's permissions or authorize a tool call.

Make the boundary concrete

Use a dedicated, read-only identity and an approved document set. Keep credentials out of prompts. Record the user, retrieved document versions and output. Do not connect a general-purpose assistant directly to programmable logic controllers or safety systems. Any later integration needs a separate engineering and security assessment.

Ask about the remote support path

Before adding AI, map who can remotely maintain the line and how that access is approved. For this pilot, require named sessions, defined expiry and an owner who can terminate them. Check that the assistant's infrastructure cannot become a new shortcut around existing plant-access controls.

Rehearse a boring failure

Disconnect the assistant. Can technicians still find the approved procedures? Can the

quality team use its established process without waiting for a model provider? The desired fallback is a practiced routine, not a newly improvised spreadsheet during a production problem.

Do not patch by surprise

The pilot is not permission to scan, update or reboot live production assets. Route plant changes through the responsible engineer, vendor constraints and the site's change process. NIST's OT guidance is the reference point for balancing security with safe, reliable operation. [\[4\]](#)

No AI suggestion should create a new route around an existing safety decision.

AGENTSSTACK PERSPECTIVE

Prove value before granting authority.

Agentsstack proposed pilot: an approved-document assistant for one maintenance team. The sequence is a starting design, not a validated delivery guarantee.

Days 1–7 / Define and baseline

Choose one equipment family and a small group of technicians. Record lookup time, the number of searches that fail and the time spent checking answers. Identify the procedure owner. Remove superseded material from the approved collection, without deleting records required for retention.

Days 8–14 / Build in isolation

Configure a read-only assistant against copied, approved documents. Create questions with known answers, unsupported questions and cases involving the wrong equipment variant. Keep test credentials separate from plant credentials.

Days 15–21 / Observe with users

Let technicians use the assistant during agreed tasks, with the existing procedure as the authority. Measure net task time after checking and correction. Capture rejected answers and why they were rejected. Limit use to the approved equipment family.

Days 22–30 / Decide, do not drift

Review outcomes with operations, quality and IT/security. Approve continued limited use only if the agreed quality, access and usability gates are met. Hold the pilot if document quality is the real constraint. Stop it if checking the output costs more time than the assistance saves.

Evidence to retain

Keep the baseline, evaluation questions, approved answers, document versions, test results and a record of configuration changes. This becomes the starting audit trail for the pilot; it is not a claim of regulatory compliance.

Hard stop conditions

Pause use after an unauthorized data retrieval, a harmful instruction presented as approved, an unexplained permission change or loss of the fallback process. Investigate before restart.

Recommended review group: plant owner · procedure owner · technician representative · IT/security owner. Name actual people before starting.

Count recovered capacity, not imaginary savings.

A worked example for a maintenance-document assistant. Every number below is an editorial assumption in Singapore dollars, not a quote or a market benchmark.

ILLUSTRATIVE CAPACITY-VALUE EXAMPLE

S\$840

Monthly net capacity value

7.1 months

Approximate setup recovery on a capacity-value basis

The assumptions

Assume 1,200 procedure lookups each month and four minutes saved per lookup, after answer checking and correction. That is 80 hours of gross task capacity. Assume only 60% can be productively redeployed: 48 hours. Value that time at a fully loaded S\$30 an hour.

What the arithmetic says

Illustrative monthly capacity value is S\$1,440. Deduct an assumed S\$600 monthly operating cost to get S\$840 of net capacity value. With an assumed S\$6,000 setup cost, the simple capacity-value recovery period is about 7.1 months.

What it does not say

This is not cash payback unless spending is actually avoided or demonstrable additional contribution is earned. Do not count the same recovered hours as both labour savings and increased output.

The downside case

At two minutes saved and the same 60% redeployment assumption, useful capacity falls to 24 hours. Its value is S\$720; after S\$600 operating cost, only S\$120 remains. The same setup amount would take 50 months to recover.

What to agree before the money moves.

The decision rule

Approve a pilot budget against an explicit learning objective. Continue only when measured benefit beats the simplest credible alternative and remains worthwhile after review effort, hosting, support, content upkeep and integration cost. The finance owner should replace every assumption with the firm's actual figures.

Three questions for a supplier

Can we export our corpus, evaluations and logs? How are user permissions enforced

before retrieval? What exactly happens when the model, connector or subscription is unavailable? Ask for a demonstration of failure handling, not only a successful answer.

Exclusions to hold the line on

Exclude downtime and scrap benefits until they are separately measured and attributable. Tax, financing and discounting are excluded from the worked example.

Original calculation: $1,200 \times 4 \div 60 \times 60\% \times \text{S\$}30 - \text{S\$}600 = \text{S\$}840$ per month.

The go/no-go decision belongs to operations, quality and security together.

AGENTSSTACK PERSPECTIVE

Leave the meeting with a decision, not a demo.

A practical checklist for the first leadership conversation.

Choose the right starting lane

Prioritize a maintenance-document assistant when approved instructions exist and lookup work is frequent. Consider packaging-image triage when representative images and quality reviewers are available. Defer autonomous scheduling or machine-control proposals until the data, operational constraints and safety responsibilities are properly assessed.

Ask operations for one baseline

Request five representative tasks and the existing procedure for each. Observe how the work is actually performed. Include a difficult case, a missing-document case and

a case where experience—not a written rule—determines the answer.

Ask security for one map

Sketch the data sources, identities, connectors and systems the assistant could reach. Name the owner of each link. If a line cannot be explained, it is not yet ready to be connected.

Ask finance for one threshold

Agree the maximum pilot cost, the minimum useful benefit and the date on which the project must be reconsidered. Decide in advance what evidence would lead the team to stop. Avoid redefining success after the results arrive.

Keep food and beverage decisions explicit

For a bottling or food plant, identify who owns label approval, lot traceability, allergen-related information and final product release. The proposed assistant

may help people find information; it must not silently inherit those responsibilities.

The question to take to your team

Which repetitive decision would become easier if the right approved information arrived faster—and could we test that without changing production authority?

An SME-sized project can still have enterprise-grade accountability.

AGENTSSTACK PERSPECTIVE

Evidence you can inspect.

Numbered references match the articles. All sources were checked on 11 September 2026.

[1] From knowing to doing: Embodied AI gets to work

Computer Weekly · Aaron Tan · August 2026

Original reporting · public synopsis. Only the public synopsis was reviewed; the full article requires registration.

[2] Agentic AI harnesses

Australian Signals Directorate / ACSC · 11 September 2026

Government guidance. New guidance at the research cut-off; not an industry adoption survey.

[3] AI RMF Profile on Trustworthy AI in Critical Infrastructure: concept note

NIST · Page updated 17 July 2026

Work in development. A developing profile and concept note; not a final regulation.

[4] SP 800-82 Rev. 3: Guide to Operational Technology (OT) Security

NIST · September 2023

Final technical guidance. Revision 3 is cited, not the later Revision 4 draft.

[5] LLM01:2025 Prompt Injection

OWASP GenAI Security Project · 2025 risk taxonomy

Community security guidance. Risk taxonomy, not a certification or proof that a particular product is secure.

EDITORIAL NOTES

This magazine does not provide legal, financial or safety certification. Publication requires a named editor's approval and a final source check. No affiliation or endorsement by cited organizations is implied.



NEXT IN THE SERIES

25 September - Financial services & fintech

Faster service, bounded agents and the controls around
customer data.

Agentsstack / AI & Cyber Weekly

One industry. Useful AI. Accountable cybersecurity.

agentsstack.com