AI agents inside the manufacturer

What they can read, what they can change, and who decides

IN BRIEF

01 Manufacturers have moved from a single assistant to hundreds of agents in a short time. One appliance maker says it has deployed more than 800 across its manufacturing, logistics and supply chain. A steelmaker is reported to have launched more than 300 in nine months. A carmaker's maintenance assistant reads equipment manuals and fault reports and suggests a fix within seconds.

02 The suppliers are moving in different directions. One large automation supplier describes agents "capable of executing entire processes without human intervention". Another offers "manual, semi-autonomous or autonomous modes", with operators able to override at any time. The appliance maker's own logistics lead puts the risk plainly: agents "can drift, they can make mistakes".

03 This paper is about the work around the line rather than on it: the reading, checking and drafting done in maintenance, quality, supply, engineering and safety. It sets out what an agent can read, how it reaches a plant's systems, what it may change without a person, and what the rules already say about machinery, records and workers. It ends with seven questions to put to any supplier, and with a place to start.

04 The position throughout is simple. An agent prepares. A named person decides. Plant data is read, never written. The agent runs on infrastructure the manufacturer controls.

IN THIS PAPER

01 Where manufacturers are today 2
02 The work nobody has packaged 3
03 Three doors, one gate 4
04 What an agent may do 6
05 What the rules already ask 6
06 A worked example: one components manufacturer 8
07 Seven questions before an agent touches a plant 9
08 Where to start 9

PART 01

Where manufacturers are today

Over 800

AI agents one appliance maker says it has deployed across manufacturing, logistics and supply chain

Over 600

suppliers whose communication one agent at the same maker manages

Over 300

specialised agents a steelmaker is reported to have launched in nine months

Seconds

for a carmaker's maintenance assistant to suggest a fix for a system problem

Agents are arriving in numbers. GE Appliances says it "has deployed more than 800 AI agents across its manufacturing, logistics, and supply chain operations", among them a Supplier Collaboration Agent that handles communication "with more than 600 suppliers" [1]. Tata Steel is reported as "launching more than 300 specialised AI agents in the past nine months", across manufacturing, maintenance, finance, procurement, human resources and customer service [4].

The first useful job is reading. BMW's Factory Genius searches "equipment manuals, quality data, internal fault reports and planning documents" and can "suggest solutions for the specific system problem within seconds". It began as a pilot in the body shop at Dingolfing [3]. That is the pattern this paper follows: an agent that reads what the plant already knows and hands it to the person who acts.

Suppliers disagree about the person. Siemens announced at AUTOMATE 2025 a move toward "autonomous agents capable of executing entire processes without human intervention" [5]. Honeywell, launching its own agents in October 2026, says "Customers can operate Forge Cognition in manual, semi-autonomous or autonomous modes" and "Operators set the level of oversight and control and can override it at any time" [6]. A buyer is being asked to choose, often before anyone has asked where the person sits.

The people running agents say they need checking. GE Appliances' vice president for logistics told a trade publication: "They can drift, they can make mistakes, and we have to respond accordingly." The same article describes the company's focus on bounded uses "where AI outputs are reviewed by humans" [2].

What the published accounts leave out. Of the manufacturing deployments we read, almost none says where the agents run; the carmaker's says only that "cloud functionality" will make the tool available worldwide [3]. We found no public incident in a plant caused by an AI agent, and we do not claim one.

On the suppliers' side, none of the agent product pages we read commits to running the agents on the customer's own premises or with no connection to the internet.

PART 02

The work nobody has packaged

Every

library template we opened acts or alerts as soon as the workflow decides

None

of them shows a step where a person approves before a ticket, a hold or a message goes out

We listed the jobs done around a plant, from the maintenance workshop to the warranty desk, and looked for each in the template libraries of the agent and workflow platforms. The libraries cover the alarm and the dashboard well. They thin out wherever a named person signs something.

WHERE TEMPLATES EXIST	WHERE NONE WAS FOUND
Asset health alerts and maintenance tickets	**A work order a planner approves before it is raised**
Shift and downtime reports	**A nonconformance report with an approver**
Inspection against thresholds	**A batch record review before quality release**
Supplier delay alerts and chasers	**A supplier corrective action request**
Schedule suggestions	**An engineering change impact review**
Questions over plant data	**Incident reports and permit-to-work checks**
	Warranty claims with a person's decision
	Audit packs with the evidence linked

Two findings matter more than the list. First, every library template we opened acts or alerts as soon as the workflow decides. One maintenance template describes a night shift where "Night alerts automatically escalate without manual verification". One supplier's published scenario for an inspection agent "Triggers rework tasks or line hold decisions when thresholds are breached", though the page adds that it is "for demonstration purposes only" [7]. None shows a step where a person approves before a ticket, a hold or a message goes out [8].

Second, the templates we opened are built from hosted services: a hosted model, a hosted chat or messaging service, a hosted ticket system. That settles where plant data and supplier terms go before anyone has asked.

The gaps sit where the reading is slow, the records are controlled, and the decision already belongs to a named person. That makes them the safer place to begin.

PART 03

Three doors, one gate

An agent is only as useful as what it can reach. In a manufacturer that means three kinds of thing.

Documents.

Work instructions, procedures, equipment manuals, specifications, the quality manual, safety data sheets. These are indexed once as knowledge sources, with each reader's permissions kept, so an agent answers with the document it relied on and says when the documents are silent.

Systems of record.

Work orders, purchase orders, nonconformances, the asset register, parts and assemblies. Most are reached through an interface the company already runs or its supplier publishes. The company gives the platform the description of that interface, chooses which operations become tools, and decides which of them need approval.

Plant data for analysis.

Line stops, equipment readings, inspection results and cost figures, held in the historian, the data warehouse or a reporting database, and read only. The control network is not one of the doors.

All three pass one gate: every tool call is checked against the company's policy, which can allow it, record it, hold it for a person, or block it.

SIX WAYS INTO A MANUFACTURER'S SYSTEMS

	ROUTE	WHAT IT REACHES	WHAT TO KNOW
01	**A file the user uploads**	A supplier's certificate, a technician's notes, a batch record, a permit	Nothing is connected. The safest start
02	**A knowledge source**	Work instructions, manuals, specifications, safety data sheets	Permissions are kept per reader
03	**An imported description of an interface**	Work orders, purchase orders, assets, nonconformances	You choose the operations. Changes wait for approval
04	**A reporting database**	Line stops, readings, inspection results, costs	Read only
05	**A server from the tool catalogue, run in your cluster**	Databases, web search, document conversion	It runs inside; check what each tool does
06	**A supplier's own tool server**	A supplier's industrial data or operations platform	A few exist; some run only in the supplier's cloud

PART 03 · CONTINUED

20
interface descriptions taken from their publishers and a public directory; the importer refused none

3,163
operations the importer found in them

1,793
of those operations change something; 1,370 read

279 of 410
operations in the two maintenance systems change something

What is published. Maintenance and plant data platforms are well described. We took 20 interface descriptions from their publishers and a public directory, among them two maintenance management systems, an industrial data platform, a time-series database, cloud device services, and the published standards for the Asset Administration Shell and the OPC UA web interface. The platform's importer refused none and found 3,163 operations [8].

What that exercise taught us. Plant interfaces are built to change things. Of the 3,163 operations, 1,793 change something and 1,370 read. The two maintenance systems alone carry 410 operations, and 279 of them change something: create a work order, close it, adjust stock, edit an asset. Every operation in the OPC UA web interface is sent as a change, even those that only read, so the method alone cannot tell a read from a write. A technician asking about a fault needs a handful of reads. Add nothing that writes to a controller at all.

What is not published. The largest planning, execution and quality suites give an importer nothing without a sign-in: the description sits behind a customer login or is generated by each installation. Some publishers split a description across many files, which are joined into one before import. Where nothing is published, the company supplies its own description, or the agent reads from the data warehouse.

Plant networks. The safe design keeps the agent on the office side of the plant network boundary, reading from the historian's copy or a reporting database, never from the controllers.

So choose the operations at import, start with the ones that read, and add nothing that changes a record until the approval step has been seen working.

PART 04

What an agent may do

One rule runs through every agent in this paper: it prepares, and a named person decides.

01	**It reads** what the person using it could read, and no more.
02	**It drafts** the work order, the report, the claim, the handover note, with each statement tied to its source.
03	**It lists** what is missing, what is out of limits, what falls outside the procedure.
04	**It stops** before anything that changes a record or reaches a supplier, a customer or a colleague. A tool that changes something waits for approval.
05	**It decides nothing about a person.** It does not allocate shifts, score operators, or name anyone as the cause of an incident.

Four lines are drawn tighter than the rest.

Machines.

No agent writes to a machine, a controller or a line. Plant data reaches it read only. An agent may propose which maintenance to bring forward, with the readings behind it; the planner decides.

Schedules and orders.

An agent may propose a revised schedule, or list the parts a job is short. It changes no schedule and places no order.

Quality release.

An agent may list a batch record's missing entries, values outside limits and steps out of order, each with the page. It signs nothing and releases nothing.

Staff.

An agent may draft a training note from a procedure. It does not assess, rank or allocate people.

PART 05

What the rules already ask

This part needs a lawyer's read before it is relied on. It reports what the texts say; it is not advice.

Machines

Learning safety functions are a separate category. The European Union's Machinery Regulation applies from 14 January 2027 [11]. It sends software that ensures safety functions to an outside conformity assessment where it has "a fully or partially self-evolving behaviour using machine learning approaches", and says those provisions "should not apply to software incapable of learning or evolving, and programmed only to execute certain automated functions" [11]. An agent that reads plant data and drafts a work order is not a safety function.

Changes to a machine must be traceable. The regulation asks that "the tracing log of the data generated in relation to an intervention and of the versions of safety software uploaded" be kept for five years [11]. One more reason to keep agents away from the controllers.

AI in machinery has new dates. Regulation (EU) 2026/1744 moved the Machinery Regulation from Section A to Section B of the AI Act's first annex, and set high-risk duties to apply from 2 December 2027 for the uses in Annex III and 2 August 2028 for those tied to Annex I [10]. What that means for a given robot cell is a question for counsel.

PART 05 · CONTINUED

Records

A controlled record shows who changed it. For records under United States drug and device rules, a system must use “secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records”, and “Record changes shall not obscure previously recorded information” [15]. An agent that drafts a record entry leaves the entry, and the signature, to a person.

AI in medicines manufacturing has a draft annex. The European Commission consulted in 2025 on a new Annex 22 to its good manufacturing practice guidelines, which “establishes requirements for the use of AI and machine learning in the manufacturing of active substances and medicinal products” and “foresees a continuous oversight of AI systems, including change control” [16]. It is a draft; the consultation closed on 7 October 2025. Read it as the direction of travel, not the rule.

Workers

The AI Act lists as high-risk any system used “to allocate tasks based on individual behaviour or personal traits or characteristics or to monitor and evaluate the performance and behaviour of persons” at work, with duties that apply from 2 December 2027 [9] [10]. A high-risk system must be built so that it “can be effectively overseen by natural persons during the period in which they are in use” [9]. An employer must “inform workers’ representatives and the affected workers” before using one at work [9]. The Act already bans AI that infers the emotions of people in the workplace, except “for medical or safety reasons” [9].

In Great Britain, the safety regulator says: “Health and safety legislation requires a risk assessment to be undertaken for uses of AI which impact on workplace health and safety and appropriate controls to be put in place to reduce risk so far as is reasonably practicable” [17].

Connected products and plant security

MARKET	WHAT WE READ	WHAT IT ASKS
European Union	Cyber Resilience Act	Makers of connected products report actively exploited vulnerabilities from 11 September 2026; the rest applies from 11 December 2027 [12]
European Union	NIS2 Directive	Lists makers of electronics, electrical equipment, machinery and vehicles as critical; they must manage the security of their networks and systems [13]
European Union	Data Act	Connected products placed on the market after 12 September 2026 must be designed so their data is available to the user by default [14]
Saudi Arabia	Operational technology cybersecurity controls (OTCC-1:2022)	Minimum cybersecurity requirements to protect industrial control systems, as “an extension to the NCA’s Essential Cybersecurity Controls” [18]

We did not open the primary text for the United Arab Emirates or for India, and found no rule in either aimed at AI in manufacturing. A company in those markets should ask its own counsel.

PART 05 · CONTINUED

What the rules have in common

01	**A person stays in charge** of anything that affects safety, a controlled record or a worker.
02	**Every change is traceable,** to a person, a time and a version.
03	**Machines and their safety functions are held to their own rules,** and an agent that reads is outside them.
04	**The company knows where the data is** and who can reach the plant network.

PART 06

A worked example: one components manufacturer

Take a made-up maker of machined and assembled components: three plants, about 1,500 people, some 400 suppliers. Nothing here is a customer's story, and every figure is illustrative; it is section 4 put to work.

01

On the line, an operator asks what torque the work instruction sets for a fastener on a new variant. The agent answers from the plant's own instruction and shows the step. Asked about a variant the instructions do not cover yet, it says so and tells the operator to ask the shift lead.

02

In the maintenance workshop, a technician asks why a spindle keeps stopping. The agent reads the manual and two years of work orders for that machine, and lists the causes the manual names, with the orders where each was found. The technician decides the repair. The agent then turns the notes into a work order; the planner submits it.

03

At goods in, a supplier sends a certificate for a batch of bar stock. The agent sets each value on the certificate against the specification and the purchase order, and lists two values the certificate leaves out. The incoming inspector holds the batch and asks the supplier.

04

In quality, an inspector finds a bore out of tolerance on 60 parts. The agent drafts the nonconformance report from the inspector's notes and photographs: what was found, against which requirement, how many parts, and the containment the procedure asks for. The quality engineer completes and approves it.

05

In engineering, a drawing change is proposed. The agent lists what it touches: the parts and assemblies, the work instructions, the open orders and the stock on hand. The change board decides.

06

On the warranty desk, the agent checks a claim against the warranty terms and the product's service history, says which condition it meets or fails, and drafts the decision note. The warranty team decides.

Every one of these reads documents or records the company already has, writes nothing to a machine, changes no record by itself, and leaves a record of who asked and what was shown.

PART 07

Seven questions before an agent touches a plant

01 **Where does it run?**
On our infrastructure, in our private cloud, or in the supplier's service? Where is the model? Can it run with no connection to the internet?

02 **What can it read?**
Does it keep each person's permissions, or read as one all-seeing account?

03 **What can it change?**
Which tools change a work order, a schedule, an order or a quality record, and does each wait for a person?

04 **Can it reach a machine?**
Show us that no tool writes to a controller, and which side of the plant network boundary the agent sits on.

05 **Who decides?**
Show us the step where a named person approves, and what the agent does when nobody has.

06 **What is recorded?**
Who asked, which tool ran, what was decided, and for how long is that kept?

07 **Can we stop it?**
One agent, one tool, or all of them, at once, by us.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and nothing on the line waits on the agent.

Questions on your own procedures

Operators and technicians ask what the plant's own instructions and manuals say and get the answer with the step cited. No production system is touched, so it can run on the first day.

Supplier certificates

The agent sets each value on a certificate against the specification and the order. The inspector accepts or rejects the material.

Nonconformance reports

The agent drafts the report from the inspector's notes. The quality engineer completes and approves it.

One line, one plant, invented data first. Then beside the manual process, measured against it.

About Katonic AI

Katonic AI builds the platform that organisations use to build, run and govern AI agents on their own infrastructure: in their own data centre, in a private or sovereign cloud, or with no connection to the internet at all. The platform is in production with governments, banks, telecom operators and utilities. Katonic is certified to ISO 27001.

KATONIC.AI/SOLUTIONS/MANUFACTURING

SOURCES

Read on 21 November 2026. Figures are the publishers' own and were not audited by us.

01 GE Appliances, "GE Appliances reinvents manufacturing operations at scale with Google Cloud's Gemini Enterprise", 22 April 2026. geappliancesco.com

02 engineering.com, "Inside GE Appliances' AI agent rollout", 28 May 2026. engineering.com

03 BMW Group press release, "Just ask Factory Genius: how AI helps maintain manufacturing equipment", 2 July 2025. newsroom.bmwgroup.com

04 CFOtech Asia, "Tata Steel expands Google Cloud AI rollout across plants", 24 April 2026. cfotech.asia

05 Siemens, on AI agents for industry announced at AUTOMATE 2025, 12 May 2025. news.siemens.com

06 Honeywell press release on the launch of Forge Cognition, 14 October 2026. honeywell.com

07 n8n workflow template 11946, on IoT incident management with technician alerts, n8n.io; Microsoft Copilot scenario library, "Manufacturing inspector agent", adoption.microsoft.com.

08 Template libraries of agent and workflow platforms, searched 21 November 2026; and 20 interface descriptions fetched from their publishers and a public directory the same day and run through the platform's importer (Katonic research).

09 Regulation (EU) 2024/1689 (AI Act), Annex III point 4(b); Article 14(1); Article 26(7); Article 5(1)(f).

10 Regulation (EU) 2026/1744, 8 July 2026, Official Journal 24 July 2026, recitals 40 and 42.

11 Regulation (EU) 2023/1230 (Machinery Regulation), Article 54; recital 55; Annex III section 1.2.1(f).

12 Regulation (EU) 2024/2847 (Cyber Resilience Act), Articles 14(1) and 71.

13 Directive (EU) 2022/2555 (NIS2), Article 21(1) and Annex II point 5.

14 Regulation (EU) 2023/2854 (Data Act), Articles 3(1) and 50.

15 United States Code of Federal Regulations, 21 CFR 11.10(e), read at law.cornell.edu.

16 European Commission, stakeholders' consultation on EudraLex Volume 4, Chapter 4, Annex 11 and new Annex 22, closed 7 October 2025. health.ec.europa.eu

17 Health and Safety Executive, "HSE's regulatory approach to artificial intelligence (AI)", 12 June 2026. press.hse.gov.uk

18 National Cybersecurity Authority (Saudi Arabia), Operational Technology Cybersecurity Controls (OTCC-1:2022) page. nca.gov.sa